

Completing Low-Rank Matrices with Corrupted Samples from Few Coefficients in General Basis

Hongyang Zhang^{†‡}, Zhouchen Lin^{‡*}, Chao Zhang^{‡*}

[†]Machine Learning Department, School of Computer Science, Carnegie Mellon University

[‡]Key Laboratory of Machine Perception (MOE), School of EECS, Peking University

^{*}Cooperative Medianet Innovation Center, Shanghai Jiao Tong University

hongyanz@cs.cmu.edu, zlin@pku.edu.cn, chzhang@cis.pku.edu.cn

Abstract

Subspace recovery from corrupted and missing data is crucial for various applications in signal processing and information theory. To complete missing values and detect column corruptions, existing robust Matrix Completion (MC) methods mostly concentrate on recovering a low-rank matrix from few corrupted coefficients w.r.t. standard basis, which, however, does not apply to more general basis, e.g., Fourier basis. In this paper, we prove that the range space of an $m \times n$ matrix with rank r can be exactly recovered from few coefficients w.r.t. *general basis*, though r and the number of corrupted samples are both as high as $O(\min\{m, n\}/\log^3(m+n))$. Our model covers previous ones as special cases, and robust MC can recover the intrinsic matrix with a higher rank. Moreover, we suggest a universal choice of the regularization parameter, which is $\lambda = 1/\sqrt{\log n}$. By our $\ell_{2,1}$ filtering algorithm, which has theoretical guarantees, we can further reduce the computational cost of our model. As an application, we also find that the solutions to extended robust Low-Rank Representation and to our extended robust MC are mutually expressible, so both our theory and algorithm can be applied to the subspace clustering problem with missing values under certain conditions. Experiments verify our theories.

Index Terms

Robust Matrix Completion, General Basis, Subspace Recovery, Outlier Detection, $\ell_{2,1}$ Filtering Algorithm

I. INTRODUCTION

We are now in an era of big and high-dimensional data. Unfortunately, due to the storage difficulty and the computational obstacle, we can measure only a few entries from the data matrix. So restoring all of the information that the data carry through the partial measurements is of great interest in data analysis. This challenging problem is also known as the Matrix Completion (MC) problem, which is highly related to the so-called recommendation system, where one tries to predict unrevealed users' preference according to the incomplete rating feedback. Admittedly, this inverse problem is ill-posed as there should be infinite number of feasible solutions. Fortunately, most of the data are structured, e.g., face [1], texture [2], and motion [3], [4], [5]. They typically lie around low-dimensional subspaces. Because the rank of data matrix corresponds to the dimensionality of subspace, recent work [6], [7], [8], [9] in convex optimization demonstrates a remarkable fact: it is possible to exactly complete an $m \times n$ matrix of rank r , if the number of randomly selected matrix elements is no less than $O((m+n)r \log^2(m+n))$.

Yet it is well known that the traditional MC model suffers from the robustness issue. It is even sensitive to minor corruptions, which commonly occur due to sensor failures and uncontrolled environments. In the recommendation system, for instance, malicious manipulation of even a single rater might drive the output of MC algorithm far from the ground truth. To resolve the issue, several efforts have been devoted to robustifying the MC model, among which robust MC [10] is the one with solid theoretical analysis. Chen et al. [10] proved that robust MC is able to exactly recover the ground truth subspace and detect the column corruptions (i.e., some entire columns are corrupted by noises), if the dimensionality of subspace is not too high and the corrupted columns are sparse compared with the input size. Most importantly, the observed expansion coefficients should be sufficient w.r.t. the standard matrix basis $\{e_i e_j^T\}_{ij}$ (please refer to Table I for explanation of notations).

However, recent advances in theoretical physics measure quantum-state entries by tomography w.r.t. the Pauli basis, which is rather different from the standard matrix one [8]. So it is not very straightforward to apply the existing theory on robust MC to such a special case. This paper tries to resolve the problem. More generally, we demonstrate the exact recoverability of an extended robust MC model in the presence of only a few coefficients w.r.t. a set of general basis, although some columns of the intrinsic matrix might be arbitrarily corrupted. By applying our $\ell_{2,1}$ filtering algorithm which has theoretical guarantees, we are able to speed up solving the model numerically. There are various applications of our results.

A. Practical Applications

In numerical analysis, instead of the standard polynomial basis $\{x^k\}_{k=0}^n$, the Legendre polynomials are widely used to represent smooth functions due to their orthogonality. Such expansions, however, are typically sensitive to perturbation: a small perturbation of the function might arbitrarily drive the fitting result far from its original. Moreover, to reduce the storage

and the computational costs, sometimes we can record only a few expansion coefficients. To complete the missing values and get the outliers removed, this paper justifies the possibility of doing so.

In digital signal processing, one usually samples the signals, e.g., voices and feature vectors, at random in the Fourier basis. However, due to sensor failure, a group of signals that we capture may be rather unreliable. To recover the intrinsic information that the signals carry and remove the outliers simultaneously, our theoretical analysis guarantees the success of robust MC w.r.t. the Fourier basis.

In quantum information theory, to obtain a maximum likelihood estimation of a quantum state of 8 ions, one typically requires hundred of thousands of measurements w.r.t. the Pauli basis, which are unaffordable because of high experimental costs. To overcome the difficulty, Gross [8] compressed the number of observations w.r.t. any basis by an MC model. However, their model is fragile to severe corruptions, which commonly occurs because of measurement errors. To robustify the model, this paper justifies the exact recoverability of robust MC w.r.t. general basis, even if the datasets are wildly corrupted.

In subspace clustering, one tries to segment the data points according to the subspaces they lie in, which can be widely applied to motion segmentation [3], [4], [5], [11], [12], face classification [1], [13], [14], [15], system identification [16], [17], [18], and image segmentation [19], [20]. Recently, it is of great interest to cluster the subspaces while the observations w.r.t. some coordinates are missing. To resolve the issue, as an application in this paper, our theorem relates robust MC to a certain subspace clustering model – the so-called extended robust Low-Rank Representation (LRR). Thus one could hope to correctly recover the structure of multiple subspaces, if robust MC is able to complete the unavailable values and remove the outlier samples at an overwhelming probability. This is guaranteed by our paper.

B. Related Work

Suppose that L_0 is an $m \times n$ data matrix of rank r whose columns are sample points, and entries are partially observed among the set \mathcal{K}_{obs} . The MC problem aims at exactly recovering L_0 , or the range space of L_0 , from the measured elements. Probably the most well-known MC model was proposed by Candès et al. [6]. To choose the lowest-rank matrix so as to fit the observed entries, the original model is formulated as

$$\min_L \text{rank}(L), \quad \text{s.t.} \quad \langle L, e_i e_j^T \rangle = \langle L_0, e_i e_j^T \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \quad (1)$$

This model, however, is untractable because problem (1) is NP-hard. Inspired by recent work in compressive sensing, Candès et al. replaced the rank in the objective function with the nuclear norm, which is the sum of singular values and is the convex envelope of rank on the unit ball of matrix operator norm. Namely,

$$\min_L \|L\|_*, \quad \text{s.t.} \quad \langle L, e_i e_j^T \rangle = \langle L_0, e_i e_j^T \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \quad (2)$$

It is worth noting that model (2) is only w.r.t. the standard matrix basis $\{e_i e_j^T\}_{ij}$. To extend the model to any basis $\{\omega_{ij}\}_{ij}$, Gross [8] proposed a more general MC model:

$$\min_L \|L\|_*, \quad \text{s.t.} \quad \langle L, \omega_{ij} \rangle = \langle L_0, \omega_{ij} \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \quad (3)$$

Models (2) and (3) both have solid theoretical guarantees: recent work [6], [7], [8], [9] showed that the models are able to exactly recover the ground truth L_0 by an overwhelming probability, if \mathcal{K}_{obs} is uniformly distributed among all sets of cardinality $O((m+n)r \log^2(m+n))$. Unfortunately, these traditional MC models suffer from the robustness issue: they are even sensitive to minor corruptions, which commonly occurs due to sensor failures, uncontrolled environments, etc.

A parallel study to the MC problem is the so-called matrix recovery, namely, recovering underlying data matrix L_0 , or the range space of L_0 , from the corrupted data matrix $M = L_0 + S_0$, where S_0 is the noise. Probably the most widely used one is Principal Component Analysis (PCA). However, PCA is fragile to outliers. Even a single but severe corruption may wildly degrade the performance of PCA. To resolve the issue, much work has been devoted to robustifying PCA [21], [22], [23], [24], [25], [26], [27], [28], [29], among which a simple yet successful model to remove column corruptions is robust PCA via Outlier Pursuit:

$$\min_{L, S} \text{rank}(L) + \lambda \|S\|_{2,1}, \quad \text{s.t.} \quad M = L + S, \quad (4)$$

and its convex relaxation

$$\min_{L, S} \|L\|_* + \lambda \|S\|_{2,1}, \quad \text{s.t.} \quad M = L + S. \quad (5)$$

Outlier Pursuit has theoretical guarantees: Xu et al. [30] and our previous work [31] proved that when the dimensionality of ground truth subspace is not too high and the column-wise corruptions are sparse compared with the sample size, Outlier Pursuit is able to recover the range space of L_0 and detect the non-zero columns of S_0 at an overwhelming probability. Nowadays, Outlier Pursuit has been widely applied to subspace clustering [32], image alignment [33], texture representation [34], etc. Unfortunately, the model cannot handle the case of missing values, which significantly limits its working range in practice.

It is worth noting that the pros and cons of above-mentioned MC and Outlier Pursuit are mutually complementary. To remedy both of their limitations, recent work [10] suggested combining the two models together, resulting in robust MC – a

model that could complete the missing values and detect the column corruptions simultaneously. Specifically, it is formulated as

$$\begin{aligned} \min_{L, S} \quad & \text{rank}(L) + \lambda \|S\|_{2,1}, \\ \text{s.t.} \quad & \langle M, e_i e_j^* \rangle = \langle L + S, e_i e_j^* \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \end{aligned} \quad (6)$$

Correspondingly, the relaxed form is

$$\begin{aligned} \min_{L, S} \quad & \|L\|_* + \lambda \|S\|_{2,1}, \\ \text{s.t.} \quad & \langle M, e_i e_j^* \rangle = \langle L + S, e_i e_j^* \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \end{aligned} \quad (7)$$

Chen et al. [10] demonstrated the recoverability of model (7), namely, if the range space of L_0 is low-dimensional, the observed entries are sufficient, and the column corruptions are sparse compared with the input size, one can hope to exactly recover the range space of L_0 and detect the corrupted samples by robust MC at an overwhelming probability. It is well reported that robust MC has been widely applied to recommendation system and medical research [10]. However, the specific basis $\{e_i e_j^T\}_{ij}$ in problem (7) limits its extensible applications to more challenging tasks, such as those discussed in Section I-A.

C. Our Contributions

In this paper, we extend robust MC to more general cases, namely, the expansion coefficients are observed w.r.t. a set of general basis. We are particularly interested in the exact recoverability of this extended model. Our contributions are as follows:

- We demonstrate that the extended robust MC model succeeds at an overwhelming probability. This result broadens the working range of traditional robust MC in three aspects: 1. the choice of basis in our model is not limited to the standard one anymore; 2. with slightly stronger yet reasonable incoherence (ambiguity) conditions, our result allows $\text{rank}(L_0)$ to be as high as $O(n/\log^3 n)$ even when the number of corruptions and observations are both constant fraction of the total input size. In comparison with the existing result which requires that $\text{rank}(L_0) = O(1)$, our analysis significantly extends the succeeding range of robust MC model; 3. we suggest that the regularization parameter be chosen as $\lambda = 1/\sqrt{\log n}$, which is universal.
- We propose a so-called $\ell_{2,1}$ filtering algorithm to reduce the computational complexity of our model. Furthermore, we establish theoretical guarantees for our algorithm, which are elegantly relevant to the incoherence of the low-rank component.
- As an application, we relate the extended robust MC model to a certain subspace clustering model – extended robust LRR. So both our theory and our algorithm on the extended robust MC can be applied to the subspace clustering problem if the extended robust MC can exactly recover the data structure.

1) Novelty of Our Analysis Technique: In the analysis of the exact recoverability of the model, we novelly divide the proof of Theorem 1 into two parts: The exact recoverability of column support and the exact recoverability of column space. We are able to attack the two problems separately thanks to the idea of expanding the objective function at the well-designed points, i.e., (\tilde{L}, \tilde{S}) for the recovery of column support and (\hat{L}, \hat{S}) for the recovery of column space, respectively (see Sections IV-B1 and IV-C1 for details). This technique enables us to decouple the randomization of \mathcal{I}_0 and Ω_{obs} , and so construct the dual variables easily by standard tools like the least squares and golfing scheme. We notice that our framework is general. It not only can be applied to the proof for easier model like Outlier Pursuit [31] (though we will sacrifice a small polylog factor for the probability of outliers), but can also hopefully simplify the proof for model with more complicated formulation. That is roughly the high-level intuition why we can handle the general basis in this paper.

In the analysis for our $\ell_{2,1}$ filtering algorithm, we take advantage of the low-rank property, namely, we recover a small-sized seed matrix first and then use the linear representation to obtain the whole desired matrix. Our analysis employs tools in recent matrix concentration literature [35] to bound the size of the seed matrix, which elegantly relates to the incoherence of the underlying matrix. This is definitely consistent with the fact that, for matrix with high incoherence, we typically need to sample more columns in order to fully observe the maximal linearly independent group (see Algorithm 1 for the procedure).

The remainder of this paper is organized as follows. Section II describes the problem setup. Section III shows our theoretical results, i.e., the exact recoverability of our model. In Section IV, we present the detailed proofs of our main results. Section V proposes a novel $\ell_{2,1}$ filtering algorithm for the extended robust MC model, and establishes theoretical guarantees for the algorithm. We show an application of our analysis to subspace clustering problem, and demonstrate the validity of our theory by experiments in Section VI. Finally, Section VII concludes the paper.

II. PROBLEM SETUP

Suppose that L_0 is an $m \times n$ data matrix of rank r , whose columns are sample points. $S_0 \in \mathbb{R}^{m \times n}$ is a noise matrix, whose column support is sparse compared with the input size n . Let $M = L_0 + S_0$. Its expansion coefficients w.r.t. a set of general basis $\{\omega_{ij}\}_{ij}$, $(i, j) \in \mathcal{K}_{obs}$, are partially observed. This paper considers the exact recovery problem as defined below.

Definition 1 (Exact Recovery Problem). *The exact recovery problem investigates whether the range space of L_0 and the column support of S_0 can be exactly recovered from randomly selected coefficients of M w.r.t. general basis, provided that some columns of M are arbitrarily corrupted.*

A similar problem was proposed in [36], [37], which recovered the whole matrix L_0 and S_0 themselves if S_0 has element-wise support. However, it is worth noting that one can only hope to recover the range space of L_0 and the column support of S_0 in Definition 1, because a corrupted column can be addition of any one vector in the range space of L_0 and another appropriate vector [30], [10], [31]. Moreover, as existing work mostly concentrates on recovering a low-rank matrix from a sampling of *matrix elements*, our exact recovery problem covers this situation as a special case.

A. Model Formulations

As our exact recovery problem defines, we study an extended robust MC model w.r.t. a set of general basis. To choose the solution L with the lowest rank, the original model is formulated as

$$\begin{aligned} \min_{L, S} \text{rank}(L) + \lambda \|S\|_{2,1}, \\ \text{s.t. } \langle M, \omega_{ij} \rangle = \langle L + S, \omega_{ij} \rangle, \quad (i, j) \in \mathcal{K}_{obs}, \end{aligned} \quad (8)$$

where \mathcal{K}_{obs} is the observation index and $\{\omega_{ij}\}_{i,j=1}^{m,n}$ is a set of ortho-normal bases such that

$$\text{Span}\{\omega_{ij}, i = 1, \dots, m\} = \text{Span}\{e_i e_j^*, i = 1, \dots, m\}, \quad \forall j. \quad (9)$$

Unfortunately, problem (8) is NP-hard because the rank function is discrete. So we replace the rank in the objective function with the nuclear norm, resulting in the relaxed formulation:

$$\begin{aligned} \min_{L, S} \|L\|_* + \lambda \|S\|_{2,1}, \\ \text{s.t. } \langle M, \omega_{ij} \rangle = \langle L + S, \omega_{ij} \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \end{aligned} \quad (10)$$

For brevity, we also rewrite it as

$$\min_{L, S} \|L\|_* + \lambda \|S\|_{2,1}, \quad \text{s.t. } \mathcal{R}(L + S) = \mathcal{R}(M), \quad (11)$$

where $\mathcal{R}(\cdot) = \sum_{i,j \in \mathcal{K}_{obs}} \langle \cdot, \omega_{ij} \rangle \omega_{ij}$ is an operator which projects a matrix onto the space $\Omega_{obs} = \text{Span}\{\omega_{ij}, i, j \in \mathcal{K}_{obs}\}$, i.e., $\mathcal{R} = \mathcal{P}_{\Omega_{obs}}$.

In this paper, we show that problem (10), or equivalently problem (11), exactly recovers the range space of L_0 and the column support of S_0 , if the rank of L_0 is no higher than $O(n/\log^3 n)$, and the number of corruptions and observations are (nearly) constant fractions of the total input size. In other words, the original problem (8) can be well approximated by the relaxed problem (10).

B. Assumptions

At first sight, it seems not always possible to successfully separate M as the low-rank term plus the column-sparse one, because there seems to not be sufficient information to avoid the identifiability issues. The identifiability issues are reflected in two aspects: the true low-rank term might be sparse and the true sparse component might be low-rank, thus we cannot hopefully identify the ground truth correctly. So we require several assumptions in order to avoid such unidentifiable cases.

1) *Incoherence Conditions on the Low-Rank Term:* As an extreme example, suppose that the low-rank term has only one non-zero entry, e.g., $e_1 e_1^*$. This matrix has a one in the top left corner and zeros elsewhere, thus being both low-rank and sparse. So it is impossible to identify this matrix as the low-rank term correctly. Moreover, we cannot expect to recover the range space of this matrix from a sampling of its entries, unless we pretty much observe all of the elements.

To resolve the issue, Gross [8] introduced μ -incoherence condition to the low-rank term $L \in \mathbb{R}^{m \times n}$ in problem (3) w.r.t. the general basis $\{\omega_{ij}\}_{i,j}$:

$$\max_{ij} \|\mathcal{P}_V \omega_{ij}\|_F^2 \leq \frac{\mu r}{n}, \quad (\text{avoid column sparsity}) \quad (12a)$$

$$\max_{ij} \|\mathcal{P}_U \omega_{ij}\|_F^2 \leq \frac{\mu r}{m}, \quad (\text{avoid row sparsity}) \quad (12b)$$

$$\max_{ij} \langle UV^*, \omega_{ij} \rangle^2 \leq \frac{\mu r}{mn}, \quad (12c)$$

where $U \Sigma V^* \in \mathbb{R}^{m \times n}$ is the skinny SVD of L . Intuitively, as discussed in [8], [37], conditions (12a), (12b), and (12c) assert that the singular vectors reasonably spread out for small μ . Because problem (3), which is a noiseless version of problem (10), requires conditions (12a), (12b), and (12c) in its theoretical guarantees [8], we will set the same incoherence conditions to analyze our model (10) as well. We argue that beyond (12a), conditions (12b) and (12c) are indispensable for the exact recovery of the target matrix in our setting. As an example, let few entries in the first row of a matrix be non-zeros while




Fig. 1. Illustration of the ambiguity condition. From the left to the right, the μ' increases and the data tend to lie in a low-dimensional subspace.

all other elements are zeros. This matrix satisfies condition (12a) but does not satisfy (12b) and (12c). In this scenario the probability of recovering its column space is not very high, as we cannot guarantee to take a sample from those uncorrupted non-zero entries, when there are a large amount of noises.

So we assume that the low-rank part \tilde{L} satisfies conditions (12a), (12b), and (12c), and the low-rank component \hat{L} satisfies condition (12a), as work [31] did (please refer to Table I for explanation of notations). Though it is more natural to assume the incoherence on L_0 , the following example shows that the incoherence of L_0 does not suffice to guarantee the success of model (10) when the rank is relatively high:

Example 1. Compute $L_0 = XY^T$ as a product of $n \times r$ i.i.d. $\mathcal{N}(0, 1)$ matrices. The column support of S_0 is sampled by Bernoulli distribution with parameter a . Let the first entry of each non-zero column of S_0 be n and all other entries be zeros. Also set the observation matrix as $\mathcal{P}_{\Omega_{obs}}(L_0 + S_0)$, where Ω_{obs} is the set of observed index selected by i.i.d. $\text{Ber}(p_0)$. We adopt $n = 10,000$, $r = 0.1n$, $p_0 = 1$, and $a = 10/n$, so there are around constant number of corrupted samples in this example. Note that, here, L_0 is incoherent fulfilling conditions (12a), (12b), and (12c), while \tilde{L} and \hat{L} are not. However, the output of algorithm falsely identifies all of the corrupted samples as the clean data. So the incoherence of L_0 cannot guarantee the exact recoverability of our model.

Imposing incoherence conditions on $\tilde{L} = L_0 + \mathcal{P}_{\mathcal{I}_0}H_L$ and $\hat{L} = L_0 + \mathcal{P}_{\mathcal{U}_0}H_L$ is not so surprising: there might be multiple solutions for the optimization model, and the low-rankness/sparseness decompositions of M are non-unique (depending on which solution we are considering). Since $\tilde{L} + \tilde{S}$ and $\hat{L} + \hat{S}$ are two eligible decompositions of M related to a fixed optimal solution pair, it is natural to consider imposing incoherence on them. Specifically, we first assume incoherence conditions (12a), (12b), and (12c) on $\tilde{L} = L_0 + \mathcal{P}_{\mathcal{I}_0}H_L$. Note that these conditions guarantee that matrix \tilde{L} cannot be sparse, so we can resolve the identifiability issue for the decomposition $M = \tilde{L} + \tilde{S}$ and hopefully recover the index \mathcal{I}_0 . After that, the ambiguity between the low rankness and the row sparseness is not an issue any more, i.e., even for row-sparse underlying matrix we can still expect to recover its column space. Here is an example to illustrate this: suppose the low rank matrix is $e_1 \mathbf{1}^*$ which has ones in the first rows and zeros elsewhere, and we have known some of the columns are corrupted by noise. Remove the outlier columns. Even we cannot fully observe the remaining entries, we can still expect to recover the column space $\text{Range}(e_1)$ since the information for the range space is sufficient to us. Therefore, we only need to impose condition (12a) on $\hat{L} = L_0 + \mathcal{P}_{\mathcal{U}_0}H_L$, which asserts that \hat{L} cannot be column-sparse.

2) *Ambiguity Conditions on Column-Sparse Term:* Analogously, the column-sparse term \hat{S} has the identification issue as well. Suppose that \hat{S} is a rank-1 matrix such that a constant fraction of the columns are zeros. This matrix is both low-rank and column-sparse, which cannot be correctly identified. To avoid this case, one needs the isotropic assumption [38], or the following ambiguity condition, on the column-sparse term \hat{S} , which is introduced by [31]:

$$\|\mathcal{B}(\hat{S})\| \leq \mu', \quad (13)$$

where μ' can be any numerical constant. Here the isotropic assumption asserts that the covariance of the noise matrix is the identity. In fact, many noise models satisfy this assumption, e.g., i.i.d. Gaussian noise. So the normalized noise vector would uniformly distribute on the surface of a unit sphere centered at the origin, thus they cannot be in a low-dimensional subspace — in other words, not low-rank. Similarly, the ambiguity condition was proposed for the same purpose [31]. Geometrically, the spectral norm stands for the length of the first principal direction (we use operator \mathcal{B} to remove the scaling factor). So condition (13) asserts that the energy for each principal direction does not differ too much, namely, the data distribute around a ball (see Figure 1), and (13) holds once the directions of non-zero columns of \hat{S} scatter sufficiently randomly. Note that the isotropic assumption implies our ambiguity condition: if the columns of \hat{S} are isotropic, $\|\mathcal{B}(\hat{S})\|$ would be a constant even though the number of column support of \hat{S} is comparable to n . Thus our ambiguity condition (13) is feasible. No matter what number of non-zero columns of \hat{S} is, the assumption guarantees matrix \hat{S} not to be low-rank.

3) *Probability Model:* Our main results assume that the column support of S_0 and the entry support of measured set \mathcal{K}_{obs} obey i.i.d. Bernoulli distribution with parameter a and parameter p_0 , respectively. Such assumptions are mild because we have no further information on the positions of outlier and measurement. More specifically, we assume that $[S_0]_{:j} = [\delta_0]_j [Z_0]_{:j}$ throughout our proof, where $[\delta_0]_j \sim \text{Ber}(p)$ determines the outlier positions and $[Z_0]_{:j}$ determines the outlier values. If an event holds with a probability at least $1 - \Theta(n^{-10})$, we say that the event happens with an overwhelming probability.

4) *Other Assumptions:* Obviously, to guarantee the exact recovery of $\text{Range}(L_0)$, the noiseless samples $\mathcal{P}_{\mathcal{I}_0^\perp} L_0$ should span the same space as that of $\text{Range}(L_0)$, i.e., $\text{Range}(L_0) = \text{Range}(\mathcal{P}_{\mathcal{I}_0^\perp} L_0)$. Otherwise, only a subspace of $\text{Range}(L_0)$ can be recovered, because the noises may be arbitrarily severe. So without loss of generality, we assume $L_0 = \mathcal{P}_{\mathcal{I}_0^\perp} L_0$, as work [30], [10] did. Moreover, the noises should be identifiable, namely, they cannot lie in the ground truth $\text{Range}(L_0)$.

C. Summary of Main Notations

In this paper, matrices are denoted by capital symbols. For matrix M , we represent $M_{:j}$ or $M^{(j)}$ as the j th column of M . We denote by M_{ij} the entry at the i th row, j th column of the matrix. For matrix operators, M^* and M^\dagger represent the conjugate transpose and the Moore-Penrose pseudo-inverse of M , respectively, and $|M|$ stands for the matrix whose (i, j) -th entry is $|M_{ij}|$.

Several norms appear in this paper, both for vector and for matrix. The only vector norm we use is $\|\cdot\|_2$, which stands for the Euclidean norm or the vector ℓ_2 norm. For matrix norm, we denote by $\|\cdot\|_*$ the nuclear norm, which stands for the sum of singular values. The matrix norm analogous to the vector ℓ_2 norm is the Frobenious norm, represented by $\|\cdot\|_F$. The pseudo-norms, $\|\cdot\|_0$ and $\|\cdot\|_{2,0}$, denote the number of non-zero entries and non-zero columns of a matrix, respectively; They are not real norms because the absolute homogeneity does not hold. The convex surrogates of $\|\cdot\|_0$ and $\|\cdot\|_{2,0}$ are matrix ℓ_1 and $\ell_{2,1}$ norms, with definitions $\|M\|_1 = \sum_{ij} |M_{ij}|$ and $\|M\|_{2,1} = \sqrt{\sum_j \|M_{:j}\|_2^2}$, respectively. The dual norms of matrix ℓ_1 and $\ell_{2,1}$ norms are ℓ_∞ and $\ell_{2,\infty}$ norms, represented by $\|M\|_\infty = \max_{ij} |M_{ij}|$ and $\|M\|_{2,\infty} = \max_j \|M_{:j}\|_2$. We also denote the operator norm of operator \mathcal{P} as $\|\mathcal{P}\| = \sup_{\|M\|_F=1} \|\mathcal{P}M\|_F$.

Our analysis involves linear spaces as well. For example, \mathcal{I} and $\text{Supp}(L)$ (similarly define \mathcal{I}_0 for L_0 , we will not restate that for the following notations) denotes the column support of matrix L . Without confusion, it forms a linear subspace. We use Ω to represent the element support of a matrix, as well as the corresponding linear subspace. The column space of a matrix is written as script \mathcal{U} , while the row space is written as script \mathcal{V} or $\text{Row}(L)$. For any space \mathcal{X} , \mathcal{X}^\perp stands for the orthogonal complement of space \mathcal{X} .

We also discuss some special matrices and spaces in our analysis. For example, (L_0, S_0) denotes the ground truth. We represent $(L^*, S^*) = (L_0 + H_L, S_0 - H_S)$ as the optimal solutions of our model, where H_L and H_S guarantee the feasibility of the solution. We are especially interested in expanding the objective function at some particular points: For the exact recovery of the column support, we focus on $(\tilde{L}, \tilde{S}) = (L_0 + \mathcal{P}_{\mathcal{I}_0} H_L, S_0 - \mathcal{P}_{\mathcal{I}_0} H_S)$; for the exact recovery of the column support, we focus on $(\tilde{L}, \tilde{S}) = (L_0 + \mathcal{P}_{\mathcal{I}_0} \mathcal{P}_{\mathcal{U}_0} H_L, S_0 - \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0} \mathcal{P}_{\mathcal{U}_0} H_L)$. Another matrix we are interested in is $\mathcal{B}(S)$, which consists of normalized non-zero columns of S and belongs to the subdifferential of $\ell_{2,1}$ norm. Similarly, the space $\mathcal{T} = \{UX^* + YV^*, \forall X, Y \in \mathbb{R}^{n \times r}\}$ is highly related to the subgradient of the nuclear norm. Namely, the subgradient of nuclear norm can be written in closed form as a term in \mathcal{T} plus a term in \mathcal{T}^\perp . The projection operator to space \mathcal{T}^\perp is denoted by $\mathcal{P}_{\mathcal{T}^\perp}$, which equals $\mathcal{P}_{\mathcal{U}^\perp} \mathcal{P}_{\mathcal{V}^\perp}$.

Table I summarizes the main notations used in this paper.

III. EXACT RECOVERABILITY OF THE MODEL

Our main results in this paper show that, surprisingly, model (11) is able to exactly recover the range space of L_0 and identify the column support of S_0 with a closed-form regularization parameter, even when only a small number of expansion coefficients are measured w.r.t. general basis and a constant fraction of columns are arbitrarily corrupted. Our theorem is as follows:

Theorem 1 (Exact Recoverability Under Bernoulli Sampling). *Any solution (L^*, S^*) to the extended robust MC (11) with $\lambda = 1/\sqrt{\log n}$ exactly recovers the column space of L_0 and the column support of S_0 with a probability at least $1 - cn^{-10}$, if the column support \mathcal{I}_0 of S_0 subjects to i.i.d. $\text{Ber}(a)$, the support \mathcal{K}_{obs} subjects to i.i.d. $\text{Ber}(p_0)$, and*

$$\text{rank}(L_0) \leq \rho_r \frac{n_{(2)}}{\mu(\log n_{(1)})^3}, \quad a \leq \rho_a \frac{n_{(2)}}{\mu n(\log n_{(1)})^3}, \quad p_0 \geq \rho_p, \quad (14)$$

where $c, \rho_r < 1, \rho_a < 1$, and $\rho_p < 1$ are all constants independent of each other, and μ is the incoherence parameter in (12a), (12b), and (12c).

Remark 1. According to [37], a recovery result under the Bernoulli model with parameter p automatically implies a corresponding result for the uniform model with parameter $\Theta(np)$ at an overwhelming probability. So conditions (14) are equivalent to

$$\text{rank}(L_0) \leq \frac{\rho_r n_{(2)}}{\mu(\log n_{(1)})^3}, \quad s \leq \frac{\rho'_s n_{(2)}}{\mu(\log n_{(1)})^3}, \quad k \geq \rho'_p n_{(1)} n_{(2)}, \quad (15)$$

where the column support \mathcal{I}_0 of S_0 is uniformly distributed among all sets of cardinality s , the support \mathcal{K}_{obs} is uniformly distributed among all sets of cardinality k , and ρ_r, ρ'_s , and ρ'_p are numerical constants.

TABLE I
SUMMARY OF MAIN NOTATIONS USED IN THE PAPER.

Notations	Meanings	Notations	Meanings
\mathcal{I}	Column Support.	Ω	Element Support.
\mathcal{I}_0	$\mathcal{I}_0 \sim \text{Ber}(a)$.	Ω_{obs}	$\Omega_{obs} \sim \text{Ber}(p_0)$.
Γ	$\Gamma = \mathcal{I}_0^\perp \cap \Omega_{obs}$	Π	$\Pi = \mathcal{I}_0 \cap \Omega_{obs}$.
m, n	Size of the data matrix M .	$n_{(1)}, n_{(2)}$	$n_{(1)} = \max\{m, n\}, n_{(2)} = \min\{m, n\}$.
$\Theta(n)$	Grows in the same order of n .	$O(n)$	Grows equal to or less than the order of n .
\otimes	Tensor product.	e_i	Vector whose i th entry is 1 and others are 0s.
Capital	A matrix.	$I, \mathbf{0}, \mathbf{1}$	The identity matrix, all-zero matrix, and all-one vector.
$M_{:j}$ or $M^{(j)}$	The j th column of matrix M .	M_{ij}	The entry at the i th row and j th column of M .
M^*	Conjugate transpose of matrix M .	M^\dagger	Moore-Penrose pseudo-inverse of matrix M .
$ M $	Matrix whose (i, j) -th entry is $ M_{ij} $.	$\ \cdot\ _2$	ℓ_2 norm for vector, $\ v\ _2 = \sqrt{\sum_i v_i^2}$.
$\ \cdot\ _*$	Nuclear norm, the sum of singular values.	$\ \cdot\ _0$	ℓ_0 norm, number of non-zero entries.
$\ \cdot\ _{2,0}$	$\ell_{2,0}$ norm, number of nonzero columns.	$\ \cdot\ _1$	ℓ_1 norm, $\ M\ _1 = \sum_{i,j} M_{ij} $.
$\ \cdot\ _{2,1}$	$\ell_{2,1}$ norm, $\ M\ _{2,1} = \sqrt{\sum_j \ M_{:j}\ _2}$.	$\ \cdot\ _{2,\infty}$	$\ell_{2,\infty}$ norm, $\ M\ _{2,\infty} = \max_j \ M_{:j}\ _2$.
$\ \cdot\ _F$	Frobenious norm, $\ M\ _F = \sqrt{\sum_{i,j} M_{ij}^2}$.	$\ \cdot\ _\infty$	Infinity norm, $\ M\ _\infty = \max_{i,j} M_{ij} $.
$\ \mathcal{P}\ $	(Matrix) operator norm.	L_0, S_0	Ground truth.
L^*, S^*	Optimal solutions, $L^* = L_0 + H_L, S^* = S_0 - H_S$.	\tilde{L}, \tilde{S}	$\tilde{L} = L_0 + \mathcal{P}_{\mathcal{I}_0} H_L, \tilde{S} = S_0 - \mathcal{P}_{\mathcal{I}_0} H_S$.
\hat{L}, \hat{S}	$\hat{L} = L_0 + \mathcal{P}_{\mathcal{U}_0} H_L, \hat{S} = S_0 - \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{U}_0} H_L$	\hat{U}, \hat{V}	Left and right singular vectors of L .
$\mathcal{U}_0, \hat{\mathcal{U}}, \mathcal{U}^*$	Column space of L_0, \hat{L}, L^* .	$\mathcal{V}_0, \hat{\mathcal{V}}, \mathcal{V}^*$	Row space of L_0, \hat{L}, L^* .
$\hat{\mathcal{T}}$	Space $\hat{\mathcal{T}} = \{\hat{U}X^* + Y\hat{V}^*, \forall X, Y \in \mathbb{R}^{n \times r}\}$.	\mathcal{X}^\perp	Orthogonal complement of the space \mathcal{X} .
$\mathcal{P}_{\hat{\mathcal{U}}}, \mathcal{P}_{\hat{\mathcal{V}}}$	$\mathcal{P}_{\hat{\mathcal{U}}}M = \hat{U}\hat{\mathcal{T}}^*M, \mathcal{P}_{\hat{\mathcal{V}}}M = M\hat{V}\hat{\mathcal{T}}^*$.	$\mathcal{P}_{\hat{\mathcal{T}}^\perp}$	$\mathcal{P}_{\hat{\mathcal{T}}^\perp}M = \mathcal{P}_{\hat{\mathcal{U}}^\perp}\mathcal{P}_{\hat{\mathcal{V}}^\perp}M$.
$\mathcal{I}_0, \hat{\mathcal{I}}, \hat{\mathcal{T}}^*$	Index of outliers of S_0, \hat{S}, S^* .	$ \mathcal{I}_0 $	Outliers number of S_0 .
$X \in \mathcal{I}$	The column support of X is a subset of \mathcal{I} .	$\mathcal{B}(\hat{S})$	$\mathcal{B}(\hat{S}) = \{H : \mathcal{P}_{\hat{\mathcal{T}}^\perp}(H) = 0; H_{:j} = \frac{\hat{S}_{:j}}{\ \hat{S}_{:j}\ _2}, j \in \hat{\mathcal{I}}\}$.
$\sim \text{Ber}(p)$	Obeys Bernoulli distribution with parameter p .	$\mathcal{N}(a, b^2)$	Gaussian distribution (mean a and variance b^2).
$\text{Row}(M)$	Row space of matrix M .	$\text{Supp}(M)$	Column support of matrix M .
$\sigma_i(M)$	The i th singular value of matrix M .	$\lambda_i(M)$	The i th eigenvalue of matrix M .
ω_{ij}	General basis.	\mathcal{R}	$\mathcal{R}(\cdot) = \sum_{ij \in \mathcal{K}_{obs}} \langle \cdot, \omega_{ij} \rangle \omega_{ij}$.

A. Comparison to Previous Results

In the traditional low-rank MC problem, one seeks to complete a low-rank matrix from only a few measurements *without corruptions*. Recently, it has been shown that a constant fraction of the entries are allowed to be missing, even if the rank of intrinsic matrix is as high as $O(n/\log^2 n)$. So compared with the result, our bound in Theorem 1 is tight up to a polylog factor. Note that the polylog gap comes from the consideration of arbitrary corruptions in our analysis. When $a = 0$, our theorem partially recovers the results of [8].

In the traditional low-rank matrix recovery problem, one tries to recover a low-rank matrix, or the range space of matrix, from *fully observed* corrupted data. To this end, our previous work [31] demonstrated that a constant fraction of the columns can be corrupted, even if the rank of intrinsic matrix is as high as $O(n/\log n)$. Compared with the result, our bound in Theorem 1 is tight up to a polylog factor as well, where the polylog gap comes from the consideration of missing values in our analysis. When $p_0 = 1$, our theorem partially recovers the results of [31].

Probably the only low-rank model that can simultaneously complete the missing values, recover the ground truth subspace, and detect the corrupted samples is robust MC [10]. As a corollary, Chen et al. [10] showed that a constant fraction of columns and entries can be corrupted and missing, respectively, if the rank of L_0 is of order $O(1)$. Compared with this, though with stronger incoherence (ambiguity) conditions, our work extends the working range of robust MC model to the rank of order $O(n/\log^3 n)$. Moreover, our results consider a set of more general basis, i.e., when $\omega_{ij} = e_i e_j^T$, our theorem partially recovers the results of [10].

Wright et al. [39] produced a certificate of optimality for (L_0, S_0) for the Compressive Principal Component Pursuit, given that (L_0, S_0) is the optimal solution for Principal Component Pursuit. There are significant differences between their work and ours: 1. Their analysis assumed that certain entries are corrupted by noise, while our paper assumes that some whole columns are noisy. In some sense, theoretical analysis on column noise is more difficult than that on Principal Component Pursuit [31]. The most distinct difference is that we cannot expect our model to exactly recover L_0 and S_0 . Rather, only the *column space* of L_0 and the *column support* of S_0 can be exactly recovered [10], [30]. 2. Wright et al.'s analysis is based on the assumption that (L_0, S_0) can be recovered by Principal Component Pursuit, while our analysis is independent of this requirement.

IV. COMPLETE PROOFS OF THEOREM 1

Theorem 1 shows the exact recoverability of our extended robust MC model w.r.t. general basis. This section is devoted to proving this result.

A. Proof Sketch

We argue that it is not very straightforward to apply the existing proofs on Robust PCA/Matrix Completion to the case of general basis, since these proofs essentially require the observed entries and the outliers to be represented under the same basis [37]. To resolve the issue, generally speaking, we novelly divide the proof of Theorem 1 into two parts: The exact recoverability of column support and the exact recoverability of column space. We are able to attack the two problems separately thanks to the idea of expanding the objective function at the well-designed points, i.e., (\tilde{L}, \tilde{S}) for the recovery of column support and (\hat{L}, \hat{S}) for the recovery of column space, respectively (see Sections IV-B1 and IV-C1 for details). This technique enables us to decouple the randomization of \mathcal{I}_0 and Ω_{obs} , and so construct the dual variables easily by standard tools like the least squares and golfing scheme. We notice that our framework is general. It not only can be applied to the proof for easier model like Outlier Pursuit [31] (though we will sacrifice a small polylog factor for the probability a of outliers), but can also hopefully simplify the proof for model with more complicated formulation, e.g., decomposing the data matrix M into more than two structural components [39]. That is roughly the high-level intuition why we can handle the general basis and improve over the previous work in this paper.

Specifically, for the exact recoverability of column support, we expand the objective function at (\tilde{L}, \tilde{S}) to establish our first class of dual conditions. Though it is standard to construct dual variables by golfing scheme, many lemmas need to be generalized in the standard setting because of the existence of both \mathcal{I}_0 and Ω_{obs} . All the preliminary work is done in Appendix A. When $p_0 = 1$ or $a = 0$, we claim that our lemmas return to the ones in [37], [10], thus being more general. The idea behind the proofs is to fix \mathcal{I}_0 first and use the randomized argument for Ω_{obs} to have a one-step result, and then allow \mathcal{I}_0 to be randomized to get our desired lemmas.

For the exact recoverability of column support, similarly, we expand the objective function at (\hat{L}, \hat{S}) to establish our second class of dual conditions. We construct the dual variables by the least squares, and prove the correctness of our construction by using generalized lemmas as well. To this end, we also utilize the ambiguity condition, which guarantees that the outlier matrix cannot be low-rank. This enables us to improve the upper bound for the rankness of the ground truth matrix from $O(1)$ to our $O(n/\log^3 n)$.

In summary, our proof proceeds in two parallel lines. The steps are as follows.

- We prove the exact recoverability of column support:
 - Section IV-B1 proves the correctness of dual condition, as shown in Lemma 1. In particular, in the proof we focus on the subgradient of objective function at (\tilde{L}, \tilde{S}) .
 - Section IV-B2 shows the construction of dual variables (17), and Section IV-B3 proves its correction in Lemma 2.
- We then prove the exact recoverability of column space:
 - Section IV-C1 proves the correctness of dual condition, as shown in Lemma 3. In particular, in the proof we focus on the subgradient of objective function at (\hat{L}, \hat{S}) .
 - Section IV-C2 shows the construction of dual variables (21), and Section IV-C3 proves its correction in Lemma 5.

B. Exact Recovery of Column Support

1) *Dual Conditions*: We first establish dual conditions for the exact recovery of the column support. The following lemma shows that once we can construct dual variables satisfying certain conditions (a.k.a. dual conditions), the column support of the outliers can be exactly recovered with a high probability by solving our robust MC model (11). Basically, the proof is to find conditions which implies that 0 belongs to the subdifferential of the objective function at the desired low-rank and column-sparse solution.

Lemma 1. Let $(L^*, S^*) = (L_0 + H_L, S'_0 - H_S)$ be any solution to the extended robust MC (11), $\tilde{L} = L_0 + \mathcal{P}_{\mathcal{I}_0} H_L$, and $\tilde{S} = S_0 - \mathcal{P}_{\mathcal{I}_0} H_S$. Assume that $\|\mathcal{P}_{\Omega_{obs}^\perp} \mathcal{P}_{\tilde{\mathcal{T}}}\| \leq 1/2$, and

$$\tilde{U}\tilde{V}^* + \tilde{W} = \lambda(\tilde{F} + \mathcal{P}_{\Omega_{obs}^\perp} \tilde{D}),$$

where $\mathcal{P}_{\tilde{\mathcal{T}}} \tilde{W} = 0$, $\|\tilde{W}\| \leq 1/2$, $\mathcal{P}_{\Omega_{obs}^\perp} \tilde{F} = 0$, $\|\tilde{F}\|_{2,\infty} \leq 1/2$, and $\|\mathcal{P}_{\Omega_{obs}^\perp} \tilde{D}\|_F \leq 1/4$. Then S^* exactly recovers the column support of S_0 , i.e., $H_L \in \mathcal{I}_0$.

Proof: We first recall that the subgradients of nuclear norm and $\ell_{2,1}$ norm are as follows:

$$\begin{aligned} \partial_{\tilde{L}} \|\tilde{L}\|_* &= \{\tilde{U}\tilde{V}^* + \tilde{Q} : \tilde{Q} \in \tilde{\mathcal{T}}^\perp, \|\tilde{Q}\| \leq 1\}, \\ \partial_{\tilde{S}} \|\tilde{S}\|_{2,1} &= \{\mathcal{B}(\tilde{S}) + \tilde{E} : \tilde{E} \in \tilde{\mathcal{I}}^\perp, \|\tilde{E}\|_{2,\infty} \leq 1\}. \end{aligned}$$

According to Lemma 7 and the feasibility of (L^*, S^*) , $\mathcal{P}_{\Omega_{obs}} H_L = \mathcal{P}_{\Omega_{obs}} H_S = H_S$. Let $\tilde{S} = S'_0 - H_S + \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L = S'_0 - \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0} H_L \in \mathcal{I}_0$. Thus the pair (\tilde{L}, \tilde{S}) is feasible to problem (11). Then we have

$$\begin{aligned} &\|L_0 + H_L\|_* + \lambda\|S'_0 - H_S\|_{2,1} \\ &\geq \|\tilde{L}\|_* + \lambda\|\tilde{S}\|_{2,1} + \langle \tilde{U}\tilde{V}^* + \tilde{Q}, \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle - \lambda\langle \mathcal{B}(\tilde{S}) + \tilde{E}, \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle. \end{aligned}$$

Now adopt \tilde{Q} such that $\langle \tilde{Q}, \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle = \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_*$ and $\langle \tilde{E}, \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle = \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_{2,1}$ ¹, and note that $\langle \mathcal{B}(\tilde{S}), \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle = 0$. So we have

$$\begin{aligned} & \|L_0 + H_L\|_* + \lambda \|S'_0 - H_S\|_{2,1} \\ & \geq \|\tilde{L}\|_* + \lambda \|\tilde{S}\|_{2,1} + \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_* + \lambda \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_{2,1} + \langle \tilde{U}\tilde{V}^*, \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle. \end{aligned}$$

Notice that

$$\begin{aligned} & |\langle \tilde{U}\tilde{V}^*, \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle| \\ & = |\langle \tilde{W} - \lambda \tilde{F} - \lambda \mathcal{P}_{\Omega_{obs}} \tilde{D}, \mathcal{P}_{\mathcal{I}_0^\perp} H_L \rangle| \\ & \leq \frac{1}{2} \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_* + \frac{\lambda}{2} \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_{2,1} + \frac{\lambda}{4} \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F. \end{aligned}$$

So we have

$$\begin{aligned} & \|L_0 + H_L\|_* + \lambda \|S'_0 - H_S\|_{2,1} \\ & \geq \|\tilde{L}\|_* + \lambda \|\tilde{S}\|_{2,1} + \frac{1}{2} \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_* + \frac{\lambda}{2} \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_{2,1} - \frac{\lambda}{4} \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F. \end{aligned}$$

Also, note that

$$\begin{aligned} & \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F \\ & \leq \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F + \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\tilde{\mathcal{T}}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F \\ & \leq \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F + \frac{1}{2} \|\mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F \\ & \leq \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F + \frac{1}{2} \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F + \frac{1}{2} \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F. \end{aligned}$$

That is

$$\|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F \leq 2 \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F + \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_F.$$

Therefore, we have

$$\begin{aligned} & \|L_0 + H_L\|_* + \lambda \|S'_0 - H_S\|_{2,1} \\ & \geq \|\tilde{L}\|_* + \lambda \|\tilde{S}\|_{2,1} + \frac{1-\lambda}{2} \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_* + \frac{\lambda}{4} \|\mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L\|_{2,1}. \end{aligned}$$

Since the pair $(L_0 + H_L, S'_0 - H_S)$ is optimal to problem (11), we have

$$\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H_L = 0 \quad \text{and} \quad \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{I}_0^\perp} H_L = 0,$$

i.e., $\mathcal{P}_{\mathcal{I}_0^\perp} H_L \in \tilde{\mathcal{T}} \cap \Omega_{obs}^\perp = \{0\}$. So $H_L \in \mathcal{I}_0$. ■

By Lemma 1, to prove the exact recovery of column support, it suffices to show a dual certificate \tilde{W} such that

$$\begin{cases} \text{(a)} & \tilde{W} \in \tilde{\mathcal{T}}^\perp, \\ \text{(b)} & \|\tilde{W}\| \leq 1/2, \\ \text{(c)} & \|\mathcal{P}_{\Omega_{obs}}^\perp (\tilde{U}\tilde{V}^* + \tilde{W})\|_F \leq \lambda/4, \\ \text{(d)} & \|\mathcal{P}_{\Omega_{obs}} (\tilde{U}\tilde{V}^* + \tilde{W})\|_{2,\infty} \leq \lambda/2. \end{cases} \quad (16)$$

2) *Certification by Golfing Scheme*: The remainder of the proofs is to construct \tilde{W} which satisfies dual conditions (16). Before introducing our construction, we assume that $\mathcal{K}_{obs} \sim \text{Ber}(p_0)$ (For brevity, we also write it as $\Omega_{obs} \sim \text{Ber}(p_0)$), or equivalently $\Omega_{obs}^\perp \sim \text{Ber}(1-p_0)$. Note that Ω_{obs} has the same distribution as that of $\Omega_1 \cup \Omega_2 \cup \dots \cup \Omega_{j_0}$, where each Ω_j is drawn from $\text{Ber}(q)$, $j_0 = \lceil \log n_{(1)} \rceil$, and q fulfills $1-p_0 = (1-q)^{j_0}$ (Note that $q = \Theta(1/\log n_{(1)})$ implies $p_0 = \Theta(1)$). We construct \tilde{W} based on such a distribution.

To construct \tilde{W} , we use the golfing scheme introduced by [8] and [37]. Let $Z_{j-1} = \mathcal{P}_{\tilde{\mathcal{T}}}(\tilde{U}\tilde{V}^* - Y_{j-1})$. We construct \tilde{W} by an inductive procedure:

$$\begin{aligned} Y_j &= Y_{j-1} + q^{-1} \mathcal{P}_{\Omega_j} Z_{j-1} = q^{-1} \sum_{k=1}^j \mathcal{P}_{\Omega_k} Z_{k-1}, \\ \tilde{W} &= \mathcal{P}_{\tilde{\mathcal{T}}^\perp} Y_{j_0}. \end{aligned} \quad (17)$$

Also, we have the inductive equation:

$$Z_j = Z_{j-1} - q^{-1} \mathcal{P}_{\tilde{\mathcal{T}}} \mathcal{P}_{\Omega_j} Z_{j-1}. \quad (18)$$

¹By the duality between the nuclear norm and the operator norm, there exists a Q such that $\langle Q, \mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H \rangle = \|\mathcal{P}_{\tilde{\mathcal{T}}^\perp} \mathcal{P}_{\mathcal{I}_0^\perp} H\|_*$ and $\|Q\| \leq 1$. Thus we take $\tilde{Q} = \mathcal{P}_{\tilde{\mathcal{T}}^\perp} Q \in \tilde{\mathcal{T}}^\perp$. It holds similarly for \tilde{E} .

3) *Proofs of Dual Conditions:* We now prove that the dual variables satisfy our dual conditions. The proof basically uses the recursiveness of the dual variables that we construct.

Lemma 2. Assume that $\Omega_{obs} \sim \text{Ber}(p_0)$ and $j_0 = \lceil \log n \rceil$. Then under the other assumptions of Theorem 1, W^L given by (17) obeys the dual conditions (16).

Proof: By Lemma 11, Lemma 12 and the inductive equation (18), when $q \geq c' \mu r \log n_{(1)} / \varepsilon^2 n_{(2)}$ for some c' , the following inequalities hold with an overwhelming probability:

$$\begin{aligned} \|Z_j\|_F &< \varepsilon^j \|Z_0\|_F = \varepsilon^j \|\tilde{U} \tilde{V}^*\|_F, \\ \max_{ab} |\langle Z_j, \omega_{ab} \rangle| &< \varepsilon^j \max_{ab} |\langle Z_0, \omega_{ab} \rangle| = \varepsilon^j \max_{ab} |\langle \tilde{U} \tilde{V}^*, \omega_{ab} \rangle|. \end{aligned}$$

Now we check the three conditions in (16).

- (a) The construction (17) implies the condition (a) holds.
- (b) It holds that

$$\begin{aligned} \|\tilde{W}\| &= \|\mathcal{P}_{\tilde{\tau}^\perp} Y_{j_0}\| \\ &\leq \sum_{k=1}^{j_0} \|q^{-1} \mathcal{P}_{\tilde{\tau}^\perp} \mathcal{P}_{\Omega_k} Z_{k-1}\| \\ &= \sum_{k=1}^{j_0} \|\mathcal{P}_{\tilde{\tau}^\perp} (q^{-1} \mathcal{P}_{\Omega_k} Z_{k-1} - Z_{k-1})\| \\ &\leq \sum_{k=1}^{j_0} \|q^{-1} \mathcal{P}_{\Omega_k} Z_{k-1} - Z_{k-1}\| \\ &\leq C_0 \sqrt{\frac{n_{(1)} \log n_{(1)}}{q}} \sum_{k=1}^{j_0} \max_{ab} |\langle Z_{k-1}, \omega_{ab} \rangle| \\ &\leq C_0 \frac{1}{1-\varepsilon} \sqrt{\frac{n_{(1)} \log n_{(1)}}{q}} \sqrt{\frac{\mu r}{mn}} \\ &= C_0 \frac{1}{1-\varepsilon} \sqrt{\frac{\rho_r}{q(\log n_{(1)})^2}} \\ &\leq \frac{1}{4}, \end{aligned}$$

where the third inequality holds due to Lemma 13 and the last inequality holds once $q \geq \Theta(1/\log n_{(1)})$.

- (c) Notice that $Y_{j_0} \in \Omega_{obs}$, i.e., $\mathcal{P}_{\Omega_{obs}^\perp} Y_{j_0} = \mathbf{0}$. Then the following inequalities follow

$$\begin{aligned} \|\mathcal{P}_{\Omega_{obs}^\perp} (\tilde{U} \tilde{V}^* + \tilde{W})\|_F &= \|\mathcal{P}_{\Omega_{obs}^\perp} (\tilde{U} \tilde{V}^* + \mathcal{P}_{\tilde{\tau}^\perp} Y_{j_0})\|_F \\ &= \|\mathcal{P}_{\Omega_{obs}^\perp} (\tilde{U} \tilde{V}^* + Y_{j_0} - \mathcal{P}_{\tilde{\tau}} Y_{j_0})\|_F \\ &= \|\mathcal{P}_{\Omega_{obs}^\perp} (\tilde{U} \tilde{V}^* - \mathcal{P}_{\tilde{\tau}} Y_{j_0})\|_F \\ &= \|\mathcal{P}_{\Omega_{obs}^\perp} Z_{j_0}\|_F \\ &\leq \varepsilon^{j_0} \sqrt{r} \quad (j_0 = \lceil \log n_{(1)} \rceil \geq \log n_{(1)}) \\ &\leq \frac{\sqrt{r}}{n_{(1)}} \quad (\varepsilon < e^{-1}) \\ &\leq \frac{\sqrt{\rho_r \mu}}{\sqrt{n_{(1)}} (\log n_{(1)})^{3/2}} \\ &\leq \frac{\lambda}{4}. \end{aligned} \tag{19}$$

- (d) We first note that $\tilde{U} \tilde{V}^* + \tilde{W} = Z_{j_0} + Y_{j_0}$. It follows from (19) that

$$\|\mathcal{P}_{\Omega_{obs}} Z_{j_0}\|_{2,\infty} \leq \|\mathcal{P}_{\Omega_{obs}} Z_{j_0}\|_F \leq \varepsilon^{j_0} \sqrt{r} \leq \frac{\lambda}{8}.$$

Moreover, we have

$$\begin{aligned}
\|\mathcal{P}_{\Omega_{obs}} Y_{j_0}\|_{2,\infty} &= \|Y_{j_0}\|_{2,\infty} \\
&\leq \sum_{k=1}^{j_0} q^{-1} \|\mathcal{P}_{\Omega_k} Z_{k-1}\|_{2,\infty} \\
&\leq q^{-1} \sqrt{m} \sum_{k=1}^{j_0} \max_{ab} |\langle Z_{k-1}, \omega_{ab} \rangle| \\
&\leq \frac{1}{q} \sqrt{\frac{\mu r}{n}} \sum_{k=1}^{j_0} \varepsilon^{k-1} \\
&\leq \frac{1}{q} \sqrt{\frac{\mu r}{n_{(2)}}} \sum_{k=1}^{j_0} \varepsilon^{k-1} \\
&\leq c \log n_{(1)} \frac{1}{(\log n_{(1)})^{3/2}} \\
&\leq \frac{\lambda}{8},
\end{aligned}$$

where the fifth inequality holds once $q \geq \Theta(1/\log n_{(1)})$. Thus $\|\mathcal{P}_{\Omega_{obs}}(\tilde{U}\tilde{V}^* + \tilde{W})\|_{2,\infty} \leq \lambda/4$. \blacksquare

C. Exact Recovery of Column Space

1) *Dual Conditions*: We then establish dual conditions for the exact recovery of the column space. The following lemma shows that if we can construct dual variables satisfying certain conditions, the column space of the underlying matrix can be exactly recovered with a high probability by solving model (11).

Lemma 3 (Dual Conditions for Exact Column Space). *Let $(L^*, S^*) = (L_0 + H_L, S'_0 - H_S)$ be any solution to the extended robust MC (11), $\hat{L} = L_0 + \mathcal{P}_{\mathcal{U}_0} H_L$, and $\hat{S} = S'_0 - \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{U}_0} H_L$. Suppose that $\hat{\mathcal{V}} \cap \Gamma^\perp = \{0\}$ and*

$$\widehat{W} = \lambda(\mathcal{B}(\hat{S}) + \hat{F}),$$

where $\widehat{W} \in \hat{\mathcal{V}}^\perp \cap \Omega_{obs}$, $\|\widehat{W}\| \leq 1/2$, $\mathcal{P}_{\Gamma^\perp} \hat{F} = 0$, and $\|\hat{F}\|_{2,\infty} \leq 1/2$. Then L^* exactly recovers the column support of L_0 , i.e., $H_L \in \mathcal{U}_0$.

Proof: We first recall that the subgradients of nuclear norm and $\ell_{2,1}$ norm are as follows:

$$\begin{aligned}
\partial_{\hat{L}} \|\hat{L}\|_* &= \{\hat{U}\hat{V}^* + \hat{Q} : \hat{Q} \in \hat{\mathcal{T}}^\perp, \|\hat{Q}\| \leq 1\}, \\
\partial_{\hat{S}} \|\hat{S}\|_{2,1} &= \{\mathcal{B}(\hat{S}) + \hat{E} : \hat{E} \in \hat{\mathcal{I}}^\perp, \|\hat{E}\|_{2,\infty} \leq 1\}.
\end{aligned}$$

By the definition of subgradient, the inequality follows

$$\begin{aligned}
&\|L_0 + H_L\|_* + \lambda\|S'_0 - H_S\|_{2,1} \\
&\geq \|\hat{L}\|_* + \lambda\|\hat{S}\|_{2,1} + \langle \hat{U}\hat{V}^* + \hat{Q}, \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle - \lambda\langle \mathcal{B}(\hat{S}) + \hat{E}, \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle \\
&\geq \|\hat{L}\|_* + \lambda\|\hat{S}\|_{2,1} + \langle \hat{U}\hat{V}^*, \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle + \langle \hat{Q}, \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle - \lambda\langle \mathcal{B}(\hat{S}), \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle - \lambda\langle \hat{E}, \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle.
\end{aligned}$$

Now adopt \hat{Q} such that $\langle \hat{Q}, \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle = \|\mathcal{P}_{\hat{\mathcal{V}}^\perp} \mathcal{P}_{\mathcal{U}_0^\perp} H_L\|_*$ and $\langle \hat{E}, \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle = -\|\mathcal{P}_{\Gamma} \mathcal{P}_{\mathcal{U}_0^\perp} H_L\|_{2,1}^2$. We have

$$\begin{aligned}
&\|L_0 + H_L\|_* + \lambda\|S'_0 - H_S\|_{2,1} \\
&\geq \|\hat{L}\|_* + \lambda\|\hat{S}\|_{2,1} + \|\mathcal{P}_{\hat{\mathcal{V}}^\perp} \mathcal{P}_{\mathcal{U}_0^\perp} H_L\|_* + \lambda\|\mathcal{P}_{\Gamma} \mathcal{P}_{\mathcal{U}_0^\perp} H_L\|_{2,1} \\
&\quad - \lambda\langle \mathcal{B}(\hat{S}), \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle.
\end{aligned}$$

Notice that

$$\begin{aligned}
|\langle -\lambda\mathcal{B}(\hat{S}), \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle| &= |\langle \lambda\hat{F} - \widehat{W}, \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle| \\
&\leq |\langle \widehat{W}, \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle| + \lambda|\langle \hat{F}, \mathcal{P}_{\mathcal{U}_0^\perp} H_L \rangle| \\
&\leq \frac{1}{2} \|\mathcal{P}_{\hat{\mathcal{V}}^\perp} \mathcal{P}_{\mathcal{U}_0^\perp} H_L\|_* + \frac{\lambda}{2} \|\mathcal{P}_{\Gamma} \mathcal{P}_{\mathcal{U}_0^\perp} H_L\|_{2,1}.
\end{aligned}$$

²By the duality between the nuclear norm and the operator norm, there exists a Q such that $\langle Q, \mathcal{P}_{\hat{\mathcal{V}}^\perp} \mathcal{P}_{\mathcal{U}_0^\perp} H \rangle = \|\mathcal{P}_{\hat{\mathcal{V}}^\perp} \mathcal{P}_{\mathcal{U}_0^\perp} H\|_*$ and $\|Q\| \leq 1$. Thus we take $\hat{Q} = \mathcal{P}_{\mathcal{U}_0^\perp} \mathcal{P}_{\hat{\mathcal{V}}^\perp} Q \in \hat{\mathcal{T}}^\perp$. It holds similarly for \hat{E} .

Hence

$$\begin{aligned}
& \|\widehat{L}\|_* + \lambda\|\widehat{S}\|_{2,1} \\
& \geq \|L_0 + H_L\|_* + \lambda\|S'_0 - H_S\|_{2,1} \\
& \geq \|\widehat{L}\|_* + \lambda\|\widehat{S}\|_{2,1} + \frac{1}{2}\|\mathcal{P}_{\widehat{\mathcal{V}}^\perp}\mathcal{P}_{\mathcal{U}_0^\perp}H_L\|_* + \frac{\lambda}{2}\|\mathcal{P}_\Gamma\mathcal{P}_{\mathcal{U}_0^\perp}H_L\|_{2,1}.
\end{aligned}$$

So $\mathcal{P}_{\mathcal{U}_0^\perp}H_L \in \widehat{\mathcal{V}} \cap \Gamma^\perp = \{0\}$, i.e., $H_L \in \mathcal{U}_0$. ■

The following lemma shows that one of the conditions in Lemma 3 holds true.

Lemma 4. *Under the assumption of Theorem 1, $\widehat{\mathcal{V}} \cap \Gamma^\perp = \{0\}$.*

Proof: We first prove $p(1-p_0)\|\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Gamma^\perp}M\|_F \leq 2\|\mathcal{P}_{\widehat{\mathcal{V}}^\perp}\mathcal{P}_{\Gamma^\perp}M\|_F$ for any matrix M . Let $M' = \mathcal{P}_{\Gamma^\perp}M$. Because $\mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}}M' + \mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}^\perp}M' = 0$, we have $\|\mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}}M'\|_F = \|\mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}^\perp}M'\|_F \leq \|\mathcal{P}_{\widehat{\mathcal{V}}^\perp}M'\|_F$. Note that

$$\begin{aligned}
& (p(1-p_0))^{-1}\|\mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}}M'\|_F \\
& = (p(1-p_0))^{-1}\langle \mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}}M', \mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}}M' \rangle \\
& = \langle \mathcal{P}_{\widehat{\mathcal{V}}}M', (p(1-p_0))^{-1}\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}}M' \rangle \\
& = \langle \mathcal{P}_{\widehat{\mathcal{V}}}M', ((p(1-p_0))^{-1}\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}} - \mathcal{P}_{\widehat{\mathcal{V}}})\mathcal{P}_{\widehat{\mathcal{V}}}M' \rangle \\
& \quad + \langle \mathcal{P}_{\widehat{\mathcal{V}}}M', \mathcal{P}_{\widehat{\mathcal{V}}}M' \rangle \\
& \geq \|\mathcal{P}_{\widehat{\mathcal{V}}}M'\|_F - \frac{1}{2}\|\mathcal{P}_{\widehat{\mathcal{V}}}M'\|_F \\
& = \frac{1}{2}\|\mathcal{P}_{\widehat{\mathcal{V}}}M'\|_F,
\end{aligned}$$

where the first inequality holds due to Corollary 2. So we have

$$\|\mathcal{P}_{\widehat{\mathcal{V}}^\perp}M'\|_F \geq \|\mathcal{P}_\Gamma\mathcal{P}_{\widehat{\mathcal{V}}}M'\|_F \geq \frac{p(1-p_0)}{2}\|\mathcal{P}_{\widehat{\mathcal{V}}}M'\|_F,$$

i.e., $p(1-p_0)\|\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Gamma^\perp}M\|_F \leq 2\|\mathcal{P}_{\widehat{\mathcal{V}}^\perp}\mathcal{P}_{\Gamma^\perp}M\|_F$.

Now let $M \in \widehat{\mathcal{V}} \cap \Gamma^\perp$. Then $\mathcal{P}_{\widehat{\mathcal{V}}^\perp}\mathcal{P}_{\Gamma^\perp}M = 0$ while $\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Gamma^\perp}M = M$. So $p(1-p_0)\|M\|_F \leq 0$, i.e., $M = 0$. Therefore, $\widehat{\mathcal{V}} \cap \Gamma^\perp = \{0\}$. ■

By Lemma 3, to prove the exact recovery of column space, it suffices to show a dual certificate \widehat{W} such that

$$\begin{cases} \widehat{W} \in \widehat{\mathcal{V}}^\perp \cap \Omega_{obs}, \\ \|\widehat{W}\| \leq 1/2, \\ \mathcal{P}_\Pi\widehat{W} = \lambda\mathcal{B}(\widehat{\mathcal{S}}), \quad \Pi = \mathcal{I}_0 \cap \Omega_{obs}, \\ \|\mathcal{P}_\Gamma\widehat{W}\|_{2,\infty} \leq \lambda/2, \quad \Gamma = \mathcal{I}_0^\perp \cap \Omega_{obs}. \end{cases} \quad (20)$$

2) *Certification by Least Squares:* The remainder of proofs is to construct \widehat{W} which satisfies the dual conditions (20). Note that $\widehat{\mathcal{I}} = \mathcal{I}_0 \sim \text{Ber}(p)$. To construct W , we consider the method of least squares, which is

$$\widehat{W} = \lambda\mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \sum_{k \geq 0} (\mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi)^k \mathcal{B}(\widehat{\mathcal{S}}), \quad (21)$$

where the Neumann series is well defined due to $\|\mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi\| < 1$. Indeed, note that $\Pi \subseteq \Omega_{obs}$. So we have the identity:

$$\begin{aligned}
& \mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi \\
& = \mathcal{P}_\Pi(\mathcal{P}_{\widehat{\mathcal{V}}} + \mathcal{P}_{\Omega_{obs}^\perp} - \mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Omega_{obs}^\perp} - \mathcal{P}_{\Omega_{obs}^\perp}\mathcal{P}_{\widehat{\mathcal{V}}} + \dots)\mathcal{P}_\Pi \\
& = \mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}}}(\mathcal{P}_{\widehat{\mathcal{V}}} + \mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Omega_{obs}^\perp}\mathcal{P}_{\widehat{\mathcal{V}}} + \dots)\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_\Pi \\
& = \mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}}}(\mathcal{P}_{\widehat{\mathcal{V}}} - \mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Omega_{obs}^\perp}\mathcal{P}_{\widehat{\mathcal{V}}})^{-1}\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_\Pi \\
& = \mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}}}(\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Omega_{obs}^\perp}\mathcal{P}_{\widehat{\mathcal{V}}})^{-1}\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_\Pi.
\end{aligned}$$

By Lemma 11 and the triangle inequality, we have that $1 - (1-p_0)^{-1}\|\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Omega_{obs}}\mathcal{P}_{\widehat{\mathcal{V}}}\| < 1/2$, i.e., $\|(\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_{\Omega_{obs}}\mathcal{P}_{\widehat{\mathcal{V}}})^{-1}\| < 2/(1-p_0)$. Therefore,

$$\begin{aligned}
\|\mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp}\|^2 & = \|\mathcal{P}_\Pi\mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi\| \\
& \leq 2(1-p_0)^{-1}\|\mathcal{P}_{\widehat{\mathcal{V}}}\mathcal{P}_\Pi\|^2 \\
& \leq 2(1-p_0)^{-1}\sigma^2 \\
& < 1,
\end{aligned} \quad (22)$$

where the second inequality holds due to Corollary 3. Note that $\mathcal{P}_\Omega \widehat{W} = \lambda \mathcal{B}(\widehat{S})$ and $\widehat{W} \in \widehat{\mathcal{V}}^\perp \cap \Omega_{obs}$. So to prove the dual conditions (20), it suffices to show that

$$\begin{cases} \text{(a)} & \|\widehat{W}\| \leq 1/2, \\ \text{(b)} & \|\mathcal{P}_\Gamma \widehat{W}\|_{2,\infty} \leq \lambda/2. \end{cases} \quad (23)$$

3) *Proofs of Dual Conditions:* We now prove that the dual variables that we construct above satisfy our dual conditions.

Lemma 5. *Under the assumptions of Theorem 1, \widehat{W} given by (21) obeys dual conditions (23).*

Proof: Let $\mathcal{H} = \sum_{k \geq 1} (\mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi)^k$. Then

$$\begin{aligned} \widehat{W} &= \lambda \mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \sum_{k \geq 0} (\mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi)^k \mathcal{B}(\widehat{S}) \\ &= \lambda \mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \mathcal{B}(\widehat{S}) + \lambda \mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \mathcal{H}(\mathcal{B}(\widehat{S})), \end{aligned} \quad (24)$$

Now we check the two conditions in (23).

(a) By the assumption, we have $\|\mathcal{B}(\widehat{S})\| \leq \mu'$. Thus the first term in (24) obeys

$$\lambda \|\mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \mathcal{B}(\widehat{S})\| \leq \lambda \|\mathcal{B}(\widehat{S})\| \leq \frac{1}{4}. \quad (25)$$

For the second term, we have

$$\lambda \|\mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \mathcal{H}(\mathcal{B}(\widehat{S}))\| \leq \lambda \|\mathcal{H}\| \|\mathcal{B}(\widehat{S})\|.$$

Then according to (22) which states that $\|\mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi\|^2 \leq 2\sigma^2/(1-p_0) \triangleq \sigma_0^2$ with high probability,

$$\|\mathcal{H}\| \leq \sum_{k \geq 1} \sigma_0^{2k} = \frac{\sigma_0^2}{1 - \sigma_0^2} \leq 1.$$

So

$$\lambda \|\mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \mathcal{H}(\mathcal{B}(\widehat{S}))\| \leq \frac{1}{4}.$$

That is

$$\|\widehat{W}\| \leq \frac{1}{2}.$$

(b) Let \mathcal{G} stand for $\mathcal{G} = \sum_{k \geq 0} (\mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{P}_\Pi)^k$. Then $\widehat{W} = \lambda \mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \mathcal{G}(\mathcal{B}(\widehat{S}))$. Notice that $\mathcal{G}(\mathcal{B}(\widehat{S})) \in \mathcal{I}_0$. Thus

$$\begin{aligned} \mathcal{P}_\Gamma \widehat{W} &= \lambda \mathcal{P}_{\mathcal{I}_0^\perp} \mathcal{P}_{\widehat{\mathcal{V}}^\perp \cap \Omega_{obs}} \mathcal{G}(\mathcal{B}(\widehat{S})) \\ &= \lambda \mathcal{P}_{\mathcal{I}_0^\perp} \mathcal{G}(\mathcal{B}(\widehat{S})) - \lambda \mathcal{P}_{\mathcal{I}_0^\perp} \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{G}(\mathcal{B}(\widehat{S})) \\ &= -\lambda \mathcal{P}_{\mathcal{I}_0^\perp} \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{G}(\mathcal{B}(\widehat{S})). \end{aligned}$$

Now denote $Q \triangleq \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{G}(\mathcal{B}(\widehat{S}))$. Note that

$$\begin{aligned} \|Q_{:j}\|^2 &= \sum_i Q_{ij}^2 = \sum_i \left\langle \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \mathcal{G}(\mathcal{B}(\widehat{S})), \omega_{ij} \right\rangle^2 \\ &= \sum_i \left\langle \mathcal{B}(\widehat{S}), \mathcal{G} \mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} (\omega_{ij}) \right\rangle^2 \\ &= \sum_i \sum_{j_0} \left\langle [\mathcal{B}(\widehat{S})]_{:j_0}, \mathcal{G} \mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} (\omega_{ij}) e_{j_0} \right\rangle^2 \\ &= \sum_{j_0} \sum_i \left((G_j e_i)^* [\mathcal{B}(\widehat{S})]_{:j_0} \right)^2 \left((G_j e_i)^* \mathcal{G} \mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} (\omega_{ij}) e_{j_0} \right)^2 \\ &\leq \sum_{j_0} \left((G_j e_{i_m})^* \mathcal{G} \mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} (\omega_{i_m j}) e_{j_0} \right)^2 \\ &= \left\| (G_j e_{i_m})^* \mathcal{G} \mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} (\omega_{i_m j}) \right\|_2^2 \\ &\leq \|\mathcal{G}\| \left\| \mathcal{P}_\Pi \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} \right\| \\ &\leq \frac{1}{4}, \quad \forall j, \end{aligned}$$

where $i_m = \arg \max_i |e_i^* \mathcal{G} \mathcal{P}_\Omega \mathcal{P}_{\widehat{\mathcal{V}} + \Omega_{obs}^\perp} (e_i e_j^*) e_{j_0}|$, G_j is a unitary matrix, and the second inequality holds because of fact (22). Thus $\|\mathcal{P}_\Gamma \widehat{W}\|_{2,\infty} = \lambda \|\mathcal{P}_{\mathcal{I}_0^\perp} Q\|_{2,\infty} \leq \lambda \|Q\|_{2,\infty} \leq \lambda/2$. The proofs are completed. ■

V. ALGORITHM

It is well known that robust MC can be efficiently solved by Alternating Direction Method of Multipliers (ADMM) [40], which is probably the most widely used method for solving nuclear norm minimization problems. In this section, we develop a faster algorithm, termed $\ell_{2,1}$ filtering algorithm, to solve the same problem.

A. $\ell_{2,1}$ Filtering Algorithm

Briefly speaking, our $\ell_{2,1}$ filtering algorithm consists of two steps: recovering the ground truth subspace from a randomly selected sub-column matrix, and then processing the remaining columns via $\ell_{2,1}$ norm based linear regression, which turns out to be a least square problem.

1) *Recovering Subspace from a Seed Matrix*: To speed up the algorithm, our strategy is to focus on a small-scaled subproblem from which we can recover the same subspace as solving the whole original problem [41]. To this end, we partition the whole matrix into two blocks. Suppose that $r = \text{rank}(L) \ll \min\{m, n\}$. We randomly sample k columns from M by i.i.d. $\text{Ber}(d/n)$ (our Theorem 3 suggests choosing d as $\Theta(r \log^3 n)$), forming a submatrix M_l (for brevity, we assume that M_l is the leftmost submatrix of M). Then we can partition M , L , and S accordingly:

$$M = [M_l, M_r], \quad S = [S_l, S_r], \quad L = [L_l, L_r].$$

To recover the desired subspace $\text{Range}(L_0)$ from M_l , we solve a small-scaled problem:

$$\min_{L_l, S_l} \|L_l\|_* + \frac{1}{\sqrt{\log k}} \|S_l\|_{2,1}, \quad \text{s.t.} \quad \mathcal{R}'(M_l) = \mathcal{R}'(L_l + S_l) \in \mathbb{R}^{m \times k}, \quad (26)$$

where $\mathcal{R}'(\cdot)$ is a linear mapping restricting $\mathcal{R}(\cdot)$ on the column index of M_l . As we will show in Section V-C, when the Bernoulli parameter d is no less than a lower bound, problem (26) exactly recovers the correct subspace $\text{Range}(L_0)$ and the column support of $[S_0]_l$ with an overwhelming probability.

2) *$\ell_{2,1}$ Filtering Step*: Since $\text{Range}(L_l) = \text{Range}(L_0)$ at an overwhelming probability, each column of L_r can be represented as the linear combinations of L_l . Namely, there exists a representation matrix $Q \in \mathbb{R}^{k \times (n-k)}$ such that

$$L_r = L_l Q.$$

Note that the part S_r should have very sparse columns, so we use the following $\ell_{2,1}$ norm based linear regression problem to explore the column supports of S_r :

$$\min_{Q, S_r} \|S_r\|_{2,1}, \quad \text{s.t.} \quad \mathcal{R}'(M_r) = \mathcal{R}'(L_l Q + S_r). \quad (27)$$

If we solve problem (27) directly by using ADMM [42], the complexity of our algorithm will be nearly the same as that of solving the whole original problem. Fortunately, we can solve (27) column-wise due to the separability of $\ell_{2,1}$ norms. Let $M_r^{(i)}$, $q^{(i)}$, and $S_r^{(i)}$ represent the i th column of M_r , Q , and S_r , respectively ($i = 1, \dots, n - k$). Then problem (27) could be decomposed into $n - k$ subproblems:

$$\min_{q^{(i)}, S_r^{(i)}} \|S_r^{(i)}\|_2, \quad \text{s.t.} \quad \mathcal{R}'_i(M_r)^{(i)} = \mathcal{R}'_i(L_l q + S_r)^{(i)} \in \mathbb{R}^m, \quad i = 1, \dots, n - k. \quad (28)$$

Equivalently,

$$\min_{q^{(i)}, \mathcal{Z}'_i(S_r^{(i)})} \|\mathcal{Z}'_i(S_r^{(i)})\|_2, \quad \text{s.t.} \quad \mathcal{Z}'_i(M_r^{(i)}) = \mathcal{Y}'_i(L_l) q^{(i)} + \mathcal{Z}'_i(S_r^{(i)}) \in \mathbb{R}^{h_i}, \quad i = 1, \dots, n - k, \quad (29)$$

where \mathcal{Z}'_i is an operator functioning on a vector which wipes out the unobserved elements, \mathcal{Y}'_i is a matrix operator which wipes out the corresponding rows of a matrix, and h_i is the number of observed elements in the i th column. As least square problems, (28) admits closed-form solutions $q^{(i)} = \mathcal{Y}'_i(L_l)^\dagger \mathcal{Z}'_i(M_r^{(i)})$, $\mathcal{Z}'_i(S_r^{(i)}) = \mathcal{Z}'_i(M_r^{(i)}) - \mathcal{Y}'_i(L_l) \mathcal{Y}'_i(L_l)^\dagger \mathcal{Z}'_i(M_r^{(i)})$, $i = 1, \dots, n - k$. If $\mathcal{Z}'_i(S_r^{(i)}) \neq \mathbf{0}$, we infer that the column $M_r^{(i)}$ is corrupted by noises.

We summarize our $\ell_{2,1}$ filtering algorithm in Algorithm 1.

B. Target Rank Estimation

As we mentioned above, our algorithm requires the rank estimation r as an input. For some specific applications, e.g., background modeling [37] and photometric stereo [43], the rank of the underlying matrix is known to us due to their physical properties. However, it is not always clear how to estimate the rank for some other cases. Here we provide a heuristic strategy for rank estimation.

Our strategy is based on the multiple trials of solving subproblem (26). Namely, starting from a small r estimation, we solve subproblem (26) by subsampling. If the optimal solution L_l^* is such that $k/\text{rank}(L_l^*) \geq \Theta(\log^3 n)$, we accept the r and output; Otherwise, we increase r by a fixed step (and so increase k) and repeat the procedure until $k/n \geq 0.5$. We require $k/n < 0.5$ because the speed advantage of our $\ell_{2,1}$ algorithm vanishes if the low-rank assumption does not hold.

Algorithm 1 $\ell_{2,1}$ Filtering Algorithm for Exact Recovery of Subspace and Support

Input: Observed data matrix $\mathcal{R}(M)$ and estimated rank r (see Section V-B).
1. Randomly sample columns from $\mathcal{R}(M) \in \mathbb{R}^{m \times n}$ by $\text{Ber}(d/n)$ to form $\mathcal{R}(M_l) \in \mathbb{R}^{m \times k}$;
2. // Line 3 recovers the subspace from a seed matrix.
3. Solve small-scaled $m \times k$ problem (26) by ADMM and obtain L_l , $\text{Range}(L_0)$, and column support of S_l ;
4. For i from 1 to $n - k$
5. Conduct QR factorization on the matrix $\mathcal{Y}'_i(L_l)$ as $\mathcal{Y}'_i(L_l) = Q_i R_i$;
6. // Line 7 implements $\ell_{2,1}$ filtering to the remaining columns.
7. Recover $\mathcal{Z}'_i(S_r^{(i)}) \in \mathbb{R}^{h_i}$ by solving (29), which is $\mathcal{Z}'_i(S_r^{(i)}) = \mathcal{Y}'_i(M_r^{(i)}) - Q_i(Q_i^* \mathcal{Y}'_i(M_r^{(i)}))$;
8. **If** $\mathcal{Y}'_i(S_r^{(i)}) \neq \mathbf{0}$
9. Output “ $M_r^{(i)}$ is an outlier”;
10. **End If**
11. End For
Output: Low-dimensional subspace $\text{Range}(L_0)$ and column support of matrix S_0 .

C. Theoretical Guarantees

In this section, we establish theoretical guarantees for our $\ell_{2,1}$ filtering algorithm. Namely, Algorithm 1 is able to exactly recover the range space of L_0 and the column support of S_0 with a high probability. To this end, we show that the two steps in Section V-A succeed at overwhelming probabilities, respectively:

- To guarantee the exact recovery of $\text{Range}(L_0)$ from the seed matrix, we prove that the sampled columns in Line 1 exactly span the desired subspace $\text{Range}(L_0)$ when the columns are restricted to the set \mathcal{I}_0^\perp , i.e., $\text{Range}(\mathcal{P}_{\mathcal{I}_0^\perp} M_l) = \text{Range}(L_0)$ (see Theorem 2); Otherwise, only a subspace of $\text{Range}(L_0)$ can be recovered by Line 3. Applying Theorem 1, we justify that Line 3 recovers the ground truth subspace from the seed matrix with an overwhelming probability (see Theorem 3).
- For $\ell_{2,1}$ filtering step, we demonstrate that, though operator \mathcal{Y}'_i randomly wipes out several rows of L_l , the columns of $\mathcal{Y}'_i(L_l)$ exactly span $\text{Range}(\mathcal{Y}'_i(L_0))$ with an overwhelming probability. So by checking whether the i th column belongs to $\text{Range}(\mathcal{Y}'_i(L_0))$, the least squares problem (29) suffices to examine whether a specific column $M_r^{(i)}$ is an outlier (see Theorem 4).

1) *Analysis for Recovering Subspace from a Seed Matrix:* To guarantee the recovery of $\text{Range}(L_0)$ by Line 3, the sampled columns in Line 1 should be informative. In other words, $\text{Range}(L_0) = \text{Range}(\mathcal{P}_{\mathcal{I}_0^\perp} M_l)$. To select the smallest number of columns in Line 1, we estimate the lower bound for the Bernoulli parameter d/n . Intuitively, this problem is highly connected to the property of $\mathcal{P}_{\mathcal{I}_0^\perp} M$. For instance, suppose that in the worst case $\mathcal{P}_{\mathcal{I}_0^\perp} M$ is a matrix whose elements in the first column are ones while all other elements equal zeros. By this time, Line 1 will select the first column (the only complete basis) at a high probability if and only if $d = n$. But for $\mathcal{P}_{\mathcal{I}_0^\perp} M$ whose elements are all equal to ones, a much smaller d suffices to guarantee the success of sampling. Thus, to identify the two cases, we involve the incoherence in our analysis.

We now estimate the smallest Bernoulli parameter d in Line 1 which ensures that $\text{Range}(L_0) \subseteq \text{Range}(M_l)$, or equivalently $\text{Range}(L_0) = \text{Range}(\mathcal{P}_{\mathcal{I}_0^\perp} M_l)$, at an overwhelming probability. The following theorem illustrates the result:

Theorem 2 (Sampling a Set of Complete Basis by Line 1). *Suppose that each column of the incoherent L_0 is sampled i.i.d. by Bernoulli distribution with parameter d/n . Let $[L_0]_l$ be the selected columns from L_0 , i.e., $[L_0]_l = \sum_j \delta_j [L_0]_{:j} e_j^*$, where $\delta_j \sim \text{Ber}(d/n)$. Then with probability at least $1 - \delta$, we have $\text{Range}([L_0]_l) = \text{Range}(L_0)$, provided that $d \geq 2\mu r \log \frac{r}{\delta}$, where μ is the incoherence parameter on the row space of matrix L_0 .*

Proof: The proof of Theorem 2 can be found in the Appendices. ■

Remark 2. *Note that a large incoherence parameter on the row space implies that slightly perturbing L_0 tremendously changes its column space. So we will need more samples in order to capture enough information about the column space of L_0 .*

To guarantee the exact recovery of desired subspace from the seed matrix, the rank r of intrinsic matrix should be low enough compared with the input size (see Theorem 1). Note that Line 1, however, selects the columns by i.i.d. $\text{Ber}(d/n)$, so that the number k of sampled columns is a random variable. Roughly, k should be around d due to the fact $\mathbb{E}(k) = d$. The following lemma implies that the magnitude of k typically has the same order as that of parameter d with an overwhelming probability.

Lemma 6. *Let n be the number of Bernoulli trials and suppose that $\Omega \sim \text{Ber}(d/n)$. Then with an overwhelming probability, $|\Omega| = \Theta(d)$, provided that $d \geq c \log n$ for a numerical constant c .*

Proof: Take a perturbation ϵ such that $d/n = m/n + \epsilon$. By scalar Chernoff bound which states that

$$\mathbb{P}(|\Omega| \leq m) \leq e^{-\epsilon^2 n^2 / 2d},$$

if taking $m = d/2$, $\epsilon = d/2n$ and $d \geq c_1 \log n$ for an appropriate constant c_1 , we have

$$\mathbb{P}(|\Omega| \leq d/2) \leq e^{-d/4} \leq n^{-10}. \quad (30)$$

In the other direction, by scalar Chernoff bound again which states that

$$\mathbb{P}(|\Omega| \geq m) \leq e^{-\epsilon^2 n^2 / 3d},$$

if taking $m = 2d$, $\epsilon = -d/n$ and $d \geq c_2 \log n$ for an appropriate constant c_2 , we obtain

$$\mathbb{P}(|\Omega| \geq 2d) \leq e^{-d/3} \leq n^{-10}. \quad (31)$$

Finally, according to (30) and (31), we conclude that $d/2 < |\Omega| < 2d$ with an overwhelming probability, provided that $d \geq c \log n$ for some constant c . ■

By Theorems 1 and 2 and Lemma 6, the following theorem justifies the success of Line 3 in Algorithm 1.

Theorem 3 (Exact Recovery of Ground Truth Subspace from Seed Matrix). *Suppose that all the conditions in Theorem 1 are fulfilled for the pair $([L_0]_l, [S_0]_l)$. Then Line 3 of Algorithm 1 exactly recovers the column space of the incoherent L_0 and the column support of $[S_0]_l$ with an overwhelming probability $1 - cn^{-10}$, provided that $d \geq C_0 \mu r \log^3 n$, where c and C_0 are numerical constants, and μ is the incoherence parameter on the row space of matrix L_0 .*

2) *Analysis for $\ell_{2,1}$ Filtering:* To justify the outlier identifiability of model (29), it suffices to show that $\text{Range}(\mathcal{Y}'_i(L_l))$ is complete, i.e., $\text{Range}(\mathcal{Y}'_i(L_l)) = \text{Range}(\mathcal{Y}'_i(L_0))$. Actually, this can be proved by the following theorem:

Theorem 4 (Outlier Identifiability of $\ell_{2,1}$ Filtering). *Suppose that each row of L_l is sampled i.i.d. by Bernoulli distribution with parameter p_0 . Let $\mathcal{Y}'_i(L_l)$ be the selected rows from matrix L_l , i.e., $\mathcal{Y}'_i(L_l) = \sum_j \delta_j L_{l_j} e_j$, where $\delta_j \sim \text{Ber}(p_0)$. Then with probability at least $1 - \delta$, we have $\text{rank}(\mathcal{Y}'_i(L_l)) = r$, or equivalently $\text{Range}(\mathcal{Y}'_i(L_l)) = \text{Range}(\mathcal{Y}'_i(L_0))$, provided that $p_0 \geq 2\mu \frac{r}{m} \log \frac{r}{\delta}$, where μ is the incoherence parameter on the column space of matrix L_0 .*

Proof: The proof is similar to that of Theorem 2, where we use a property that $\mu(L_l) = \mu(L_0)$ since $\text{Range}(L_l) = \text{Range}(L_0)$, by Theorem 2. ■

It is worth noting that, when the matrix is fully observed, model (29) exactly identifies the outliers even without Theorem 4.

D. Complexity Analysis

In this section, we consider the time complexity of our randomized $\ell_{2,1}$ filtering algorithm. We analyze our algorithm in the case where $d = \Theta(r \log^3 n)$. In Algorithm 1, Line 1 requires $O(n)$ time. For Line 3 which recovers $m \times cr(\log n)^3$ seed matrix, this step requires $O(r^2 m \log^6 n)$ time. Line 5 requires at most $6r^2 m \log^3 n$ time due to the QR factorization [44], and Line 7 needs $(2r + 1)m$ time due to matrix-matrix multiplication. Thus the overall complexity of our $\ell_{2,1}$ filtering algorithm is at most $O(r^2 m \log^6 n) + (2r + 1)mn + 6r^2 mn \log^3 n \approx 6r^2 mn \log^3 n$. As ADMM algorithm requires $O(mn \min\{m, n\})$ time to run for our model due to SVD or matrix-matrix multiplication in every iteration, and require many iterations in order to converge, our algorithm is significantly faster than the state-of-the-art methods.

VI. APPLICATIONS AND EXPERIMENTS

As we discuss in Section I, our model and algorithm have various applications. To show that, this section first relates our model to the subspace clustering task with missing values, and then demonstrates the validity of our theory and applications by synthetic and real experiments.

A. Applications to Robust Subspace Clustering with Missing Values

Subspace clustering aims at clustering data according to the subspaces they lie in. It is well known that many datasets, e.g., face [1] and motion [3], [4], [5], can be well separated by their different subspaces. So subspace clustering has been successfully applied to face recognition [1], motion segmentation [45], etc.

Probably one of the most effective subspace clustering models is robust LRR [46], [47]. Suppose that the data matrix M contains columns that are from a union of independent subspaces with outliers. The idea of robust LRR is to self-express the data, namely, using the clean data themselves as the dictionary, and then find the representation matrix with the lowest rank. Mathematically, it is formulated as

$$\min_{Z, L, S} \|Z\|_* + \lambda \|S\|_{2,1}, \quad \text{s.t.} \quad L = LZ, \quad M = L + S. \quad (32)$$

After obtaining the optimal solution Z^* , we can apply spectral clustering algorithms, such as Normalized Cut, to cluster each data point according to the subspaces they lie in.

Although robust LRR (32) has been widely applied to many computer vision tasks [46], [47], it cannot handle missing values, i.e., only a few entries of M are observed. Such a situation commonly occurs because of sensor failures, uncontrolled environments, etc. To resolve the issue, in this paper we extend robust LRR by slightly modifying the second constraint:

$$\min_{Z, L, S} \|Z\|_* + \lambda \|S\|_{2,1}, \quad \text{s.t.} \quad L = LZ, \quad \langle M, \omega_{ij} \rangle = \langle L + S, \omega_{ij} \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \quad (33)$$

A similar model has been proposed by Shi et al. [48], which is

$$\min_{Z, D, S} \|Z\|_* + \lambda \|S\|_{2,1}, \quad \text{s.t.} \quad D = DZ + S, \quad \langle M, e_i e_j^* \rangle = \langle D, e_i e_j^* \rangle, \quad (i, j) \in \mathcal{K}_{obs}. \quad (34)$$

However, there are two main differences between their model and ours: 1. Our model does not require ω_{ij} to be standard basis, thus being more general; 2. Unlike (34), we use clean data as the dictionary to represent themselves. Such a modification robustifies the model significantly, as discussed in [41], [47].

The extended robust LRR (33) is NP-hard due to its non-convexity, which incurs great difficulty in efficient solution. As an application of this paper, we show that the solutions to (33) and to (8) are mutually expressible in closed forms:

Claim 1. *The pair $(L^*(L^*)^\dagger, L^*, S^*)$ is optimal to the extended robust LRR problem (33), if (L^*, S^*) is a solution to the extended robust MC problem (8). Conversely, suppose that (Z^*, L^*, S^*) is a solution to the extended robust LRR problem (33), then (L^*, S^*) will be optimal to the extended robust MC problem (8).*

Proof: The proof can be found in the Appendices. ■

Using relaxed form (10) to well approximate original problem (8) according to Theorem 1, and then applying Claim 1 to obtain a solution to the extended robust LRR problem model (33), we are able to robustly cluster subspaces even though a constant fraction of values are unobserved. This is true once the conditions in Theorem 1 can be satisfied:

- The low-rankness condition holds if the sum of the subspaces is low-dimensional (less than $O(n/\log^3 n)$);
- The incoherence condition holds if the number of the subspaces is not too large (less than an absolute constant) [49].

The computational cost can be further cut by applying our $\ell_{2,1}$ filtering approach, i.e., Algorithm 1.

Remark 3. *Intuitively, Claim 1 is equivalent to a two-step procedure: First completing the data matrix and identifying the outlier by extended robust MC, and then clustering the data by LRR.*

B. Simulations and Experiments

In this section, we conduct a series of experiments to demonstrate the validity of our theorems, and show possible applications of our model and algorithm.

1) *Validity of Regularization Parameter:* We first verify the validity of our regularization parameter $\lambda = 1/\sqrt{\log n}$ by simulations. The toy data are designed as follows. We compute $L_0 = XY^T$ as a product of two $n \times r$ i.i.d. $\mathcal{N}(0, 1)$ matrices. The non-zero columns of S_0 are sampled by Bernoulli distribution with parameter a , whose entries obey i.i.d. $\mathcal{N}(0, 1)$. Finally, we construct our observation matrix as $\mathcal{P}_{\Omega_{obs}}(L_0 + S_0)$, where Ω_{obs} is the observed index selected by i.i.d. $\text{Ber}(p_0)$. We solve model (11) to obtain an optimal solution (L^*, S^*) , and then compare it with (L_0, S_0) . The distance between the range spaces of L^* and L_0 is defined by $\|\mathcal{P}_{\mathcal{U}^*} - \mathcal{P}_{\mathcal{U}_0}\|_F$ and the distance between the column supports of S^* and S_0 is given by the Hamming distance. The experiment is run by 10 times and we report the average outputs. Table II illustrates that our choice of the regularization parameter enables model (11) to exactly recover the range space of L_0 and the column support of S_0 at a high probability.

TABLE II
EXACT RECOVERY ON PROBLEMS WITH DIFFERENT SIZES. HERE $\text{RANK}(L_0) = 0.05n$, $a = 0.1$, $p_0 = 0.8$, AND $\lambda = 1/\sqrt{\log n}$.

n	$\text{dist}(\text{Range}(L^*), \text{Range}(L_0))$	$\text{dist}(\mathcal{I}^*, \mathcal{I}_0)$
100	1.09×10^{-14}	0
200	1.70×10^{-14}	0
500	4.02×10^{-14}	0
1,000	6.10×10^{-13}	0

Theorem 1 shows that the exact recoverability of model (11) is independent of the magnitudes of noises. To verify this, Table III records the differences between the ground truth (L_0, S_0) and the output (L^*, S^*) of model (11) under varying noise magnitudes $\mathcal{N}(0, 1/n)$, $\mathcal{N}(0, 1)$, and $\mathcal{N}(0, n)$. It seems that our model always succeeds, no matter what magnitudes the noises are.

TABLE III

EXACT RECOVERY ON PROBLEMS WITH DIFFERENT NOISE MAGNITUDES. HERE $n = 200$, $\text{RANK}(L_0) = 0.05n$, $a = 0.1$, $p_0 = 0.8$, AND $\lambda = 1/\sqrt{\log n}$.

Magnitude	$\text{dist}(\text{Range}(L^*), \text{Range}(L_0))$	$\text{dist}(\mathcal{I}^*, \mathcal{I}_0)$
$\mathcal{N}(0, 1/n)$	1.98×10^{-14}	0
$\mathcal{N}(0, 1)$	1.50×10^{-14}	0
$\mathcal{N}(0, n)$	3.20×10^{-14}	0




Fig. 2. Exact recovery of the extended robust MC on random problems of varying sizes. The white region represents the exact recovery in 10 experiments, and black region denotes the failures in all of the experiments.

2) *Exact Recovery from Varying Fractions of Corruptions and Observations:* We then test the exact recoverability of our model under varying fractions of corruptions and observations. The data are generated as the above-mentioned experiments, where the data size $n = 200$. We repeat the experiments by decreasing the number of observations. Each simulation is run by 10 times, and Figure 2 plots the fraction of correct recoveries: white region represents the exact recovery in 10 experiments, and black region denotes the failures in all of the experiments. It seems that model (11) succeeds even when the rank of intrinsic matrix is comparable to $O(n)$, which is consistent with our forecasted order $O(n/\log^3 n)$. But with the decreasing number of observations, the working range of model (11) shrinks.

3) *Speed Advantage of $\ell_{2,1}$ Filtering Algorithm:* To test the speed advantage of our $\ell_{2,1}$ filtering algorithm, we compare the running time of ADMM and our filtering Algorithm 1 on the synthetic data. The data are generated as the above-mentioned simulations, where we change one variable among the set (n, r, p_0, a) each time and fix others. Table IV lists the CPU times, the distance between $\text{Range}(L^*)$ and $\text{Range}(L_0)$, and the Hamming distance between \mathcal{I}^* and \mathcal{I}_0 by the two algorithms. It is easy to see that our $\ell_{2,1}$ filtering approach is significantly faster than ADMM under a comparable precision.

TABLE IV

COMPARISON OF THE SPEED BETWEEN ADMM AND OUR $\ell_{2,1}$ FILTERING ALGORITHM UNDER VARYING PARAMETER SETTINGS.

Parameter (n, r, p_0, a)	Method	Time (s)	$\text{dist}(\text{Range}(L^*), \text{Range}(L_0))$	$\text{dist}(\mathcal{I}^*, \mathcal{I}_0)$
(1,000, 1, 0.95, 0.1)	ADMM	102.89	2.34×10^{-9}	0
	$\ell_{2,1}$ Filtering	5.86	5.41×10^{-9}	0
(2,000, 1, 0.95, 0.1)	ADMM	587.53	6.03×10^{-9}	0
	$\ell_{2,1}$ Filtering	35.49	9.91×10^{-9}	0
(1,000, 10, 0.95, 0.1)	ADMM	104.12	2.71×10^{-8}	0
	$\ell_{2,1}$ Filtering	25.75	4.06×10^{-8}	0
(1,000, 1, 0.8, 0.1)	ADMM	100.10	2.76×10^{-9}	0
	$\ell_{2,1}$ Filtering	4.33	6.17×10^{-9}	0
(1,000, 1, 0.95, 0.2)	ADMM	92.95	4.01×10^{-9}	0
	$\ell_{2,1}$ Filtering	5.09	8.34×10^{-9}	0

TABLE V
CLUSTERING ACCURACIES OF OUR ALGORITHM ON THE FIRST 5 SEQUENCES IN HOPKINS 155 DATABASE, WHERE THERE ARE 5% MISSING ENTRIES.

#Sequence	Motion Number	MC+SSC	MC+LRR	Robust MC+SSC	Robust MC+LRR (Ours)
#1	2	60.48%	80.00%	82.38%	98.10%
#2	3	67.43%	79.46%	73.24%	91.91%
#3	2	75.16%	98.00%	83.01%	98.04%
#4	2	73.09%	67.55%	86.54%	99.21%
#5	2	75.69%	82.41%	75.69%	92.59%

4) *Applications to Subspace Clustering with Missing Coefficients*: To apply our model to the subspace clustering tasks with a fraction of missing values, we conduct experiments on the real Hopkins 155 database³. The Hopkins 155 database consists of 155 sequences, each of which contains multiple key points drawn from two or three motion objects. Because the motion trajectory of each rigid body lies in a single subspace, so we are able to cluster the points according to the subspaces they lie in. To make the problem more challenging, we randomly corrupt 5% columns and remove 5% observed coefficients. Table V lists the clustering accuracies of our algorithm on the first 5 sequences in comparison with other approaches [50]. We can see that our approach always achieves high clustering accuracy, even though we cannot observe all of the data values. In addition, the experiments show that the Robust MC based methods are better than MC based methods. So our model is more robust.

VII. CONCLUSIONS

In this paper, we investigate the theory, the algorithm, and the applications of our extended robust MC model. In particular, we study the exact recoverability of our model from few observed coefficients w.r.t. general basis, which partially covers the existing results as special cases. With slightly stronger incoherence (ambiguity) conditions, we are able to push the upper bound on the allowed rank from $O(1)$ to $O(n/\log^3 n)$, even when there are around a constant fraction of unobserved coefficients and column corruptions, where n is the sample size. We further suggest a universal choice of the regularization parameter, which is $\lambda = 1/\sqrt{\log n}$. This result waives the necessity of tuning regularization parameter, so it significantly extends the working range of robust MC. Moreover, we propose $\ell_{2,1}$ filtering algorithm so as to speed up solving our model numerically, and establish corresponding theoretical guarantees. As an application, we also relate our model to the subspace clustering tasks with missing values so that our theory and algorithm can be immediately applied to the subspace segmentation problem. Our experiments on the synthetic and real data testify to our theories.

VIII. ACKNOWLEDGEMENTS

Zhouchen Lin is supported by National Basic Research Program (973 Program) of China (grant no. 2015CB352502), National Natural Science Foundation (NSF) of China (grant nos. 61272341 and 61231002), and Microsoft Research Asia Collaborative Research Program. Chao Zhang is supported in part by 973 Program of China (grant nos. 2015CB352303 and 2011CB302400) and in part by NSF of China (grant nos. 61071156 and 61131003).

REFERENCES

- [1] G. Liu, Z. Lin, S. Yan, J. Sun, and Y. Ma, "Robust recovery of subspace structures by low-rank representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 171–184, 2013.
- [2] Y. Ma, H. Derksen, W. Hong, and J. Wright, "Segmentation of multivariate mixed data via lossy data coding and compression," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 9, pp. 1546–1562, 2007.
- [3] W. Gear, "Multibody grouping from motion images," *International Journal of Computer Vision*, vol. 29, no. 2, pp. 133–150, 1998.
- [4] J. Yan and M. Pollefeys, "A general framework for motion segmentation: Independent, articulated, rigid, non-rigid, degenerate and nondegenerate," in *European Conference on Computer Vision*, vol. 3954, 2006, pp. 94–106.
- [5] S. Rao, R. Tron, R. Vidal, and Y. Ma, "Motion segmentation in the presence of outlying, incomplete, or corrupted trajectories," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 10, pp. 1832–1845, 2010.
- [6] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
- [7] E. J. Candès and T. Tao, "The power of convex relaxation: Near-optimal matrix completion," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2053–2080, 2010.
- [8] D. Gross, "Recovering low-rank matrices from few coefficients in any basis," *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1548–1566, 2011.
- [9] B. Recht, "A simpler approach to matrix completion," *Journal of Machine Learning Research*, vol. 12, pp. 3413–3430, 2011.
- [10] Y. Chen, H. Xu, C. Caramanis, and S. Sanghavi, "Robust matrix completion and corrupted columns," in *International Conference on Machine Learning*, 2011, pp. 873–880.
- [11] J. Costeira and T. Kanade, "A multibody factorization method for independently moving objects," *International Journal of Computer Vision*, vol. 29, no. 3, pp. 159–179, 1998.
- [12] R. Vidal and R. Hartley, "Motion segmentation with missing data using PowerFactorization and GPCA," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2004, pp. 85–105.

³<http://www.vision.jhu.edu/data/hopkins155>

- [13] J. Ho, M. Yang, J. Lim, K. Lee, and D. Kriegman, "Clustering appearances of objects under varying illumination conditions," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2003, pp. 313–320.
- [14] R. Vidal, Y. Ma, and S. Sastry, "Generalized principal component analysis (GPCA)," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 12, pp. 1945–1959, 2005.
- [15] G. Liu and S. Yan, "Latent low-rank representation for subspace segmentation and feature extraction," in *IEEE International Conference on Computer Vision*, 2011, pp. 1615–1622.
- [16] R. Vidal, S. Soatto, Y. Ma, and S. Sastry, "An algebraic geometric approach to the identification of a class of linear hybrid systems," in *IEEE International Conference on Decision and Control*, 2003, pp. 167–172.
- [17] C. Zhang and R. Bitmead, "Subspace system identification for training-based MIMO channel estimation," *Automatica*, vol. 41, no. 9, pp. 1623–1632, 2005.
- [18] S. Paoletti, A. Juloski, G. Ferrari-Trecate, and R. Vidal, "Identification of hybrid systems—a tutorial," *European Journal of Control*, vol. 13, no. 2–3, pp. 242–260, 2007.
- [19] A. Yang, J. Wright, Y. Ma, and S. Sastry, "Unsupervised segmentation of natural images via lossy data compression," *Computer Vision and Image Understanding*, vol. 110, no. 2, pp. 212–225, 2008.
- [20] B. Cheng, G. Liu, J. Wang, H. Li, and S. Yan, "Multi-task low-rank affinity pursuit for image segmentation," in *IEEE International Conference on Computer Vision*, 2011, pp. 2439–2446.
- [21] R. Gnanadesikan and J. Kettenring, "Robust estimates, residuals, and outlier detection with multiresponse data," *Biometrics*, vol. 28, no. 1, pp. 81–124, 1972.
- [22] P. J. Huber, *Robust statistics*. Springer, 2011.
- [23] M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [24] F. De La Torre and M. J. Black, "A framework for robust subspace learning," *International Journal of Computer Vision*, vol. 54, no. 1, pp. 117–142, 2003.
- [25] Q. Ke and T. Kanade, "Robust ℓ_1 -norm factorization in the presence of outliers and missing data by alternative convex programming," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2005, pp. 739–746.
- [26] M. McCoy, J. A. Tropp *et al.*, "Two proposals for robust PCA using semidefinite programming," *Electronic Journal of Statistics*, vol. 5, pp. 1123–1160, 2011.
- [27] T. Zhang and G. Lerman, "A novel m-estimator for robust PCA," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 749–808, 2014.
- [28] G. Lerman, M. B. McCoy, J. A. Tropp, and T. Zhang, "Robust computation of linear models by convex relaxation," *Foundations of Computational Mathematics*, pp. 1–48, 2014.
- [29] M. Hardt and A. Moitra, "Algorithms and hardness for robust subspace recovery," *arXiv preprint: 1211.1041*, 2012.
- [30] H. Xu, C. Caramanis, and S. Sanghavi, "Robust PCA via outlier pursuit," *IEEE Transaction on Information Theory*, vol. 58, no. 5, pp. 3047–3064, 2012.
- [31] H. Zhang, Z. Lin, C. Zhang, and E. Chang, "Exact recoverability of robust PCA via outlier pursuit with tight recovery bounds," in *AAAI Conference on Artificial Intelligence*, 2015, pp. 3143–3149.
- [32] H. Zhang, Z. Lin, C. Zhang, and J. Gao, "Robust latent low rank representation for subspace clustering," *Neurocomputing*, vol. 145, pp. 369–373, 2014.
- [33] Y. Peng, A. Ganesh, J. Wright, W. Xu, and Y. Ma, "RASL: Robust alignment by sparse and low-rank decomposition for linearly correlated images," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 763–770.
- [34] Z. Zhang, A. Ganesh, X. Liang, and Y. Ma, "TILT: Transform-invariant low-rank textures," *International Journal of Computer Vision*, vol. 99, no. 1, pp. 1–24, 2012.
- [35] J. A. Tropp, "An introduction to matrix concentration inequality," *arXiv preprint: 1501.01571*, 2015.
- [36] H. Zhang, Z. Lin, and C. Zhang, "A counterexample for the validity of using nuclear norm as a convex surrogate of rank," in *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, vol. 8189, 2013, pp. 226–241.
- [37] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of the ACM*, vol. 58, no. 3, p. 11, 2011.
- [38] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," *arXiv preprint: 1011.3027*, 2010.
- [39] J. Wright, A. Ganesh, K. Min, and Y. Ma, "Compressive principal component pursuit," *Information and Inference*, vol. 2, no. 1, pp. 32–68, 2013.
- [40] Z. Lin, R. Liu, and Z. Su, "Linearized alternating direction method with adaptive penalty for low-rank representation," in *Advances in Neural Information Processing Systems*, 2011, pp. 612–620.
- [41] H. Zhang, Z. Lin, C. Zhang, and J. Gao, "Relations among some low rank subspace recovery models," *Neural Computation*, vol. 27, pp. 1915–1950, 2015.
- [42] R. Liu, Z. Lin, F. D. L. Torre, and Z. Su, "Fixed-rank representation for unsupervised visual learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 598–605.
- [43] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma, "Robust photometric stereo via low-rank matrix completion and recovery," in *Asian Conference on Computer Vision*, 2011, pp. 703–717.
- [44] G. Golub and C. Van Loan, *Matrix computations*. Johns Hopkins University Press, 2012, vol. 3.
- [45] E. Elhamifar and R. Vidal, "Sparse subspace clustering," in *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2, 2009, pp. 2790–2797.
- [46] P. Favaro, R. Vidal, and A. Ravichandran, "A closed form solution to robust subspace estimation and clustering," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 1801–1807.
- [47] S. Wei and Z. Lin, "Analysis and improvement of low rank representation for subspace segmentation," *arXiv preprint: 1107.1561*, 2010.
- [48] J. Shi, W. Yang, L. Yong, and X. Zheng, "Low-rank representation for incomplete data," *Mathematical Problems in Engineering*, vol. 2014, 2014.
- [49] G. Liu and P. Li, "Recovery of coherent data via low-rank dictionary pursuit," in *Advances in Neural Information Processing Systems*, 2014, pp. 1206–1214.
- [50] C. Yang, D. Robinson, and R. Vidal, "Sparse subspace clustering with missing entries," in *International Conference on Machine Learning*, 2015.
- [51] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [52] A. Gittens and J. A. Tropp, "Tail bounds for all eigenvalues of a sum of random matrices," *arXiv preprint: 1104.4513*, 2011.

APPENDIX

A. Preliminary Lemmas

We present several preliminary lemmas here which are critical for our proofs. For those readers who are interested in the main body of the proofs, please refer to Sections IV-B and IV-C directly.

Lemma 7. *The optimal solution (L^*, S^*) to the extended robust MC (11) satisfies $S^* \in \Omega_{obs}$.*

Proof: Suppose that $S^* \notin \Omega_{obs}$. We have $\|L^*\|_* + \|\mathcal{P}_{\Omega_{obs}} S^*\|_{2,1} < \|L^*\|_* + \|S^*\|_{2,1}$. Also, notice that the pair $(L^*, \mathcal{P}_{\Omega_{obs}} S^*)$ is feasible to problem (11). Thus we have a contradiction to the optimality of (L^*, S^*) . ■

Lemma 8 (Elimination Lemma on Observed Elements). *Suppose that any solution (L^*, S^*) to the extended robust MC (11) with observation set \mathcal{K}_{obs} exactly recovers the column space of L_0 and the column support of S_0 , i.e., $\text{Range}(L^*) = \text{Range}(L_0)$ and $\{j : S_{:,j}^* \notin \text{Range}(L^*)\} = \mathcal{I}_0$. Then any solution (L^*, S'^*) to (11) with observation set \mathcal{K}'_{obs} succeeds as well, where $\mathcal{K}_{obs} \subseteq \mathcal{K}'_{obs}$.*

Proof: The conclusion holds because the constraints in problem (11) with observation set \mathcal{K}'_{obs} are stronger than the constraints in problem (11) with observation set \mathcal{K}_{obs} . ■

Lemma 9 (Elimination Lemma on Column Support). *Suppose that any solution (L^*, S^*) to the extended robust MC (11) with input $\mathcal{R}(M) = \mathcal{R}(L^*) + \mathcal{R}(S^*)$ exactly recovers the column space of L_0 and the column support of S_0 , i.e., $\text{Range}(L^*) = \text{Range}(L_0)$ and $\{j : S_{:,j}^* \notin \text{Range}(L^*)\} = \mathcal{I}_0$. Then any solution (L^*, S'^*) to (11) with input $\mathcal{R}(M') = \mathcal{R}(L^*) + \mathcal{R}\mathcal{P}_{\mathcal{I}}(S^*)$ succeeds as well, where $\mathcal{I} \subseteq \mathcal{I}^* = \mathcal{I}_0$.*

Proof: Since (L^*, S'^*) is the solution of (11) with input matrix $\mathcal{P}_{\Omega_{obs}} M'$, we have

$$\|L'^*\|_* + \lambda \|S'^*\|_{2,1} \leq \|L^*\|_* + \lambda \|\mathcal{P}_{\mathcal{I}} S^*\|_{2,1}.$$

Therefore

$$\begin{aligned} & \|L'^*\|_* + \lambda \|S'^* + \mathcal{P}_{\mathcal{I}^\perp \cap \mathcal{I}_0} S^*\|_{2,1} \\ & \leq \|L'^*\|_* + \lambda \|S'^*\|_{2,1} + \lambda \|\mathcal{P}_{\mathcal{I}^\perp \cap \mathcal{I}_0} S^*\|_{2,1} \\ & \leq \|L^*\|_* + \lambda \|\mathcal{P}_{\mathcal{I}} S^*\|_{2,1} + \lambda \|\mathcal{P}_{\mathcal{I}^\perp \cap \mathcal{I}_0} S^*\|_{2,1} \\ & = \|L^*\|_* + \lambda \|S^*\|_{2,1}. \end{aligned}$$

Note that

$$\mathcal{R}(L'^* + S'^* + \mathcal{P}_{\mathcal{I}^\perp \cap \mathcal{I}_0} S^*) = \mathcal{R}(M' + \mathcal{P}_{\mathcal{I}^\perp \cap \mathcal{I}_0} S^*) = \mathcal{R}(M).$$

Thus $(L'^*, S'^* + \mathcal{P}_{\mathcal{I}^\perp \cap \mathcal{I}_0} S^*)$ is optimal to problem with input $\mathcal{P}_{\Omega_{obs}} M$ and by assumption we have

$$\text{Range}(L'^*) = \text{Range}(L^*) = \text{Range}(L_0),$$

$$\{j : [S'^* + \mathcal{P}_{\mathcal{I}^\perp \cap \mathcal{I}_0} S^*]_{:,j} \notin \text{Range}(L_0)\} = \text{Supp}(S_0).$$

The second equation implies $\mathcal{I} \subseteq \{j : S_{:,j}'^* \notin \text{Range}(L_0)\}$. Suppose that $\mathcal{I} \neq \{j : S_{:,j}'^* \notin \text{Range}(L_0)\}$. Then there exists an index k such that $S_{:,k}'^* \notin \text{Range}(L_0)$ and $k \notin \mathcal{I}$, i.e., $M'_{:,k} = L_{:,k}^* \in \text{Range}(L_0)$. Note that $L_{:,j}^* \in \text{Range}(L_0)$. Thus $S_{:,k}'^* \in \text{Range}(L_0)$ and we have a contradiction. Thus $\mathcal{I} = \{j : S_{:,j}'^* \notin \text{Range}(L_0)\} = \{j : S_{:,j}'^* \notin \text{Range}(L^*)\}$ and the algorithm succeeds. ■

Lemma 10 (Matrix (Operator) Bernstein Inequality [35]). *Let $X_i \in \mathbb{R}^{m \times n}$, $i = 1, \dots, s$, be independent, zero-mean, matrix-valued random variables. Assume that $V, L \in \mathbb{R}$ are such that $\max\{\|\sum_{i=1}^s \mathbb{E}[X_i X_i^*]\|, \|\sum_{i=1}^s \mathbb{E}[X_i^* X_i]\|\} \leq M$ and $\|X_i\| \leq L$. Then*

$$\mathbb{P}\left[\left\|\sum_{i=1}^s X_i\right\| > t\right] \leq (m+n) \exp\left(-\frac{3t^2}{8M}\right)$$

for $t \leq M/L$, and

$$\mathbb{P}\left[\left\|\sum_{i=1}^s X_i\right\| > t\right] \leq (m+n) \exp\left(-\frac{3t}{8L}\right)$$

for $t > M/L$.

Lemma 9 shows that the success of algorithm is monotone on $|\mathcal{I}_0|$. Thus by standard arguments in [37], [51], and [7], any guarantee proved for the Bernoulli distribution equivalently holds for the uniform distribution.

Lemma 11. *For any $\mathcal{K} \sim \text{Ber}(p)$, with high probability,*

$$\|\mathcal{P}_{\tilde{\mathcal{T}}} - p^{-1} \mathcal{P}_{\tilde{\mathcal{T}}} \mathcal{R}' \mathcal{P}_{\tilde{\mathcal{T}}}\| < \varepsilon \quad \text{and} \quad \|\mathcal{P}_{\hat{\mathcal{V}}} - p^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\hat{\mathcal{V}}}\| < \varepsilon,$$

provided that $p \geq C_0 \varepsilon^{-2} (\mu r \log n_{(1)}) / n_{(2)}$ for some numerical constant $C_0 > 0$, where $\mathcal{R}'(\cdot) = \sum_{ij \in \mathcal{K}} \langle \cdot, \omega_{ij} \rangle \omega_{ij}$.

Proof: The proof is in Appendix D. ■

Corollary 1 ([37]). *Assume that $\mathcal{K}_{obs} \sim \text{Ber}(p_0)$. Then with an overwhelming probability, $\|\mathcal{P}_{\Omega_{obs}^\perp} \mathcal{P}_{\tilde{\mathcal{T}}}\|^2 \leq \varepsilon + 1 - p_0$, provided that $p_0 \geq C_0 \varepsilon^{-2} (\mu r \log n) / n$ for some numerical constant $C_0 > 0$.*

Lemma 12. Suppose that $Z \in \tilde{\mathcal{T}}$ and $\mathcal{K} \sim \text{Ber}(p)$. Let $\mathcal{R}'(\cdot) = \sum_{ij \in \mathcal{K}} \langle \cdot, \omega_{ij} \rangle \omega_{ij}$. Then with high probability

$$\max_{ab} |\langle Z - p^{-1} \mathcal{P}_{\tilde{\mathcal{T}}} \mathcal{R}' Z, \omega_{ab} \rangle| < \varepsilon \max_{ab} |\langle Z, \omega_{ab} \rangle|,$$

provided that $p \geq C_0 \varepsilon^{-2} (\mu r \log n_{(1)}) / n_{(2)}$ for some numerical constant $C_0 > 0$.

Proof: The proof is in Appendix E. ■

Lemma 13. Suppose that Z is a fixed matrix and $\mathcal{K} \sim \text{Ber}(p)$. Let $\mathcal{R}'(\cdot) = \sum_{ij \in \mathcal{K}} \langle \cdot, \omega_{ij} \rangle \omega_{ij}$. Then with high probability

$$\|Z - p^{-1} \mathcal{R}' Z\| < C'_0 \sqrt{\frac{n_{(1)} \log n_{(1)}}{p}} \max_{ij} |\langle Z, \omega_{ij} \rangle|,$$

provided that $p \geq C'_0 (\mu \log n_{(1)}) / n_{(1)}$ for some small numerical constant $C'_0 > 0$.

Proof: The proof is in Appendix F. ■

Lemma 14. Let \mathcal{R}' be the projection operator onto space $\Omega = \text{Span}\{\omega_{ij}, i, j \in \mathcal{K}\}$ with any \mathcal{K} , the space $\mathcal{I} = \text{Span}\{\omega_{ij}, j \in \mathcal{J}\}$, and $\Psi = \mathcal{I} \cap \Omega$. Let $\mathcal{J} \sim \text{Ber}(a)$. Then with high probability

$$\|a^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\mathcal{I}} \mathcal{R}' \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\hat{\mathcal{V}}}\| = \|a^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Psi} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\hat{\mathcal{V}}}\| < \varepsilon,$$

provided that $a \geq C_0 \varepsilon^{-2} (\mu r \log n_{(1)}) / n$ for some numerical constant $C_0 > 0$.

Proof: The proof is in Appendix G. ■

Corollary 2. Assume that $\Psi = \mathcal{I} \cap \Omega$. Then for any $\mathcal{I} \sim \text{Ber}(a)$ and $\Omega \sim \text{Ber}(p)$, with high probability

$$\|(pa)^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Psi} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}}\| < (p^{-1} + 1) \varepsilon,$$

provided that $a, p \geq C_0 \varepsilon^{-2} (\mu r \log n_{(1)}) / n$ for some numerical constant $C_0 > 0$.

Proof: By Lemma 11 and Lemma 14, we have

$$\|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega} \mathcal{P}_{\hat{\mathcal{V}}} - p \mathcal{P}_{\hat{\mathcal{V}}}\| < p \varepsilon,$$

and

$$\|a^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Gamma} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega} \mathcal{P}_{\hat{\mathcal{V}}}\| < \varepsilon.$$

So by triangle inequality, we have

$$\begin{aligned} & \|a^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Gamma} \mathcal{P}_{\hat{\mathcal{V}}} - p \mathcal{P}_{\hat{\mathcal{V}}}\| \\ & \leq \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega} \mathcal{P}_{\hat{\mathcal{V}}} - p \mathcal{P}_{\hat{\mathcal{V}}}\| + \|a^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Gamma} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega} \mathcal{P}_{\hat{\mathcal{V}}}\| \\ & < (p + 1) \varepsilon. \end{aligned}$$

That is

$$\|(pa)^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Gamma} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}}\| < (p^{-1} + 1) \varepsilon. ■$$

Corollary 3. Let $\Pi = \mathcal{I}_0 \cap \Omega_{obs}$, where $\mathcal{I}_0 \sim \text{Ber}(p_1)$. Then with an overwhelming probability $\|\mathcal{P}_{\Pi} \mathcal{P}_{\hat{\mathcal{V}}}\|^2 \leq (1 - p_1) \varepsilon + p_1$, provided that $1 - p_1 \geq C_0 \varepsilon^{-2} (\mu r \log n_{(1)}) / n$ for some numerical constant $C_0 > 0$.

Proof: Let $\Gamma = \mathcal{I}_0^\perp \cap \Omega_{obs}$. Note that $\mathcal{I}_0^\perp \sim \text{Ber}(1 - p_1)$. By Lemma 14, we have $\|(1 - p_1)^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Gamma} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| < \varepsilon$, or equivalently

$$\begin{aligned} & \|(1 - p_1)^{-1} \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Gamma} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| \\ & = (1 - p_1)^{-1} \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Gamma} \mathcal{P}_{\hat{\mathcal{V}}} - (1 - p_1) \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| \\ & = (1 - p_1)^{-1} \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}} - \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{(\mathcal{I}_0^\perp \cap \Omega_{obs})} \mathcal{P}_{\hat{\mathcal{V}}} - p_1 \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| \\ & = (1 - p_1)^{-1} \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{(\mathcal{I}_0 \cap \Omega_{obs})} \mathcal{P}_{\hat{\mathcal{V}}} - p_1 \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| \\ & = (1 - p_1)^{-1} \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Pi} \mathcal{P}_{\hat{\mathcal{V}}} - p_1 \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| \\ & < \varepsilon. \end{aligned}$$

Therefore, by the triangle inequality

$$\|\mathcal{P}_{\Pi} \mathcal{P}_{\hat{\mathcal{V}}}\|^2 = \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Pi} \mathcal{P}_{\hat{\mathcal{V}}}\| \leq \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Pi} \mathcal{P}_{\hat{\mathcal{V}}} - p_1 \mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| + p_1 \|\mathcal{P}_{\hat{\mathcal{V}}} \mathcal{P}_{\Omega_{obs}} \mathcal{P}_{\hat{\mathcal{V}}}\| \leq (1 - p_1) \varepsilon + p_1. ■$$

B. Proofs of Theorem 2

To prove Theorem 2, the following matrix Chernoff Bound is invoked in our proof:

Theorem 5 (Matrix Chernoff Bound [52]). *Consider a finite sequence $\{X_k\} \in \mathbb{R}^{d \times d}$ of independent, random, Hermitian matrices. Assume that*

$$0 \leq \lambda_{\min}(X_k) \leq \lambda_{\max}(X_k) \leq L.$$

Define $Y = \sum_k X_k$, and μ_r as the r th largest eigenvalue of the expectation $\mathbb{E}Y$, i.e., $\mu_r = \lambda_r(\mathbb{E}Y)$. Then

$$\begin{aligned} \mathbb{P}\{\lambda_r(Y) > (1 - \epsilon)\mu_r\} &\geq 1 - r \left[\frac{e^{-\epsilon}}{(1 - \epsilon)^{1-\epsilon}} \right]^{\frac{\mu_r}{L}} \\ &\geq 1 - re^{-\frac{\mu_r \epsilon^2}{2L}}, \end{aligned}$$

for $\epsilon \in [0, 1)$.

Lemma 15. *Let $X = U\Sigma V^T$ be the skinny SVD of matrix X . For any set of coordinates Ω and any matrix $X \in \mathbb{R}^{m \times n}$, we have $\text{rank}(X_{\Omega,:}) = \text{rank}(U_{\Omega,:})$ and $\text{rank}(X_{:, \Omega}) = \text{rank}(V_{:, \Omega})$.*

Proof: On one hand,

$$X_{\Omega,:} = I_{\Omega}X = I_{\Omega}U\Sigma V^T = U_{\Omega,:}\Sigma V^T.$$

So $\text{rank}(X_{\Omega,:}) \leq \text{rank}(U_{\Omega,:})$. On the other hand, we have

$$X_{\Omega,:}V\Sigma^{-1} = U_{\Omega,:}.$$

Thus $\text{rank}(U_{\Omega,:}) \leq \text{rank}(X_{\Omega,:})$. So $\text{rank}(X_{\Omega,:}) = \text{rank}(U_{\Omega,:})$.

The second part of the argument can be proved similarly. Indeed, $X_{:, \Omega} = U\Sigma V^T I_{\Omega} = U\Sigma[V^T]_{:, \Omega}$ and $\Sigma^{-1}U^T X_{:, \Omega} = [V^T]_{:, \Omega}$. So $\text{rank}(X_{:, \Omega}) = \text{rank}([V^T]_{:, \Omega}) = \text{rank}(V_{:, \Omega})$, as desired. \blacksquare

Now we are ready to prove Theorem 2.

Proof: We investigate the smallest sampling parameter d such that the sampled columns from $L_0 = \mathcal{P}_{\mathcal{I}_0^\perp} M$ exactly span $\text{Range}(L_0)$ with an overwhelming probability.

Denote by $L_0 = U\Sigma V^T$ the skinny SVD of L_0 . Let $X = \sum_i \delta_i [V^T]_{:, i} e_i^*$ be the random sampling of columns from matrix V^T , where $\delta_i \sim \text{Ber}(d/n)$. Define a positive semi-definite matrix

$$Y = XX^* = \sum_{i=1}^n \delta_i [V^T]_{:, i} [V^T]_{:, i}^*.$$

Obviously, $\sigma_r(X)^2 = \lambda_r(Y)$. To invoke the matrix Chernoff bound, we need to estimate L and μ_r in Theorem 5. Specifically, since

$$\begin{aligned} \mathbb{E}Y &= \sum_{i=1}^n \mathbb{E}\delta_i [V^T]_{:, i} [V^T]_{:, i}^* \\ &= \frac{d}{n} \sum_{i=1}^n [V^T]_{:, i} [V^T]_{:, i}^* \\ &= \frac{d}{n} [V^T] [V^T]^*, \end{aligned}$$

we have $\mu_r = \lambda_r(\mathbb{E}Y) = d/n$. Furthermore, we also have

$$\lambda_{\max}(\delta_i [V^T]_{:, i} [V^T]_{:, i}^*) = \|\delta_i [V^T]_{:, i}\|_2^2 \leq \max_i \|V_i\|_{2, \infty}^2 = \frac{\mu r}{n} \triangleq L.$$

By matrix Chernoff bound,

$$\begin{aligned} \mathbb{P}\{\sigma_r(X) > 0\} &\geq 1 - re^{-\frac{\mu r}{2L}} \\ &= 1 - re^{-d/(2\mu r)} \\ &\geq 1 - \delta, \end{aligned}$$

we obtain

$$d \geq 2\mu r \log\left(\frac{r}{\delta}\right).$$

Note that $\sigma_r(X) > 0$ implies that $\text{rank}([V^T]_{:, i}) = \text{rank}([L_0]_{:, i}) = r$, where the equality holds due to Lemma 15. Also, $\text{Range}([L_0]_{:, i}) \subseteq \text{Range}(L_0)$. Thus $\text{Range}([L_0]_{:, i}) = \text{Range}(L_0)$. \blacksquare

C. Proofs of Claim 1

To prove Theorem 1, the following proposition is crucial throughout our proof.

Proposition 1. *The solution to the optimization problem:*

$$\min_Z \|Z\|_*, \quad \text{s.t. } L = LZ, \quad (35)$$

is unique and given by $Z^* = V_L V_L^T$, where $U_L \Sigma_L V_L^T$ is the skinny SVD of L .

Proof: We only prove the former part of the theorem. The proofs for the latter part of the theorem are similar. Suppose that (L^*, S^*) is a solution to problem (8), while $(L^*(L^*)^\dagger, L^*, S^*)$ is not optimal to problem (33). So there exists an optimal solution to (33), termed (Z_*, L_*, S_*) , which is strictly better than $(L^*(L^*)^\dagger, L^*, S^*)$. Namely,

$$\begin{aligned} \|Z_*\|_* + \lambda \|S_*\|_{2,1} &< \|L^*(L^*)^\dagger\|_* + \lambda \|S^*\|_{2,1}, \\ L_* &= L_* Z_*, \quad \mathcal{R}(M) = \mathcal{R}(L_* + S_*). \end{aligned}$$

Fixing L and S as L_* and S_* in (33), respectively, and by Proposition 1, we have

$$\begin{aligned} \|Z_*\|_* + \lambda \|S_*\|_{2,1} &= \|V_{L_*} V_{L_*}^T\|_* + \lambda \|S_*\|_{2,1} \\ &= \text{rank}(L_*) + \lambda \|S_*\|_{2,1}. \end{aligned}$$

Furthermore, by the property of Moore-Penrose pseudo-inverse,

$$\|L^*(L^*)^\dagger\|_* + \lambda \|S^*\|_{2,1} = \text{rank}(L^*) + \lambda \|S^*\|_{2,1}.$$

Thus

$$\begin{aligned} \text{rank}(L_*) + \lambda \|S_*\|_{2,1} &< \text{rank}(L^*) + \lambda \|S^*\|_{2,1}, \\ L_* &= L_* Z_*, \quad \mathcal{R}(M) = \mathcal{R}(L_* + S_*), \end{aligned}$$

which is contradictory to the optimality of (L^*, S^*) to problem (8). So $(L^*(L^*)^\dagger, L^*, S^*)$ is optimal to problem (33). \blacksquare

D. Proofs of Lemma 11

Now we are prepared to prove Lemma 11.

Proof: For any matrix X , we have

$$\mathcal{P}_{\mathcal{X}} X = \sum_{ij} \langle \mathcal{P}_{\mathcal{X}} X, \omega_{ij} \rangle \omega_{ij},$$

where \mathcal{X} is $\widehat{\mathcal{V}}$ or $\widetilde{\mathcal{T}}$. Thus $\mathcal{R}' \mathcal{P}_{\mathcal{X}} X = \sum_{ij} \kappa_{ij} \langle \mathcal{P}_{\mathcal{X}} X, \omega_{ij} \rangle \omega_{ij}$, where κ_{ij} s are i.i.d. Bernoulli variables with parameter p . Then

$$\mathcal{P}_{\mathcal{X}} \mathcal{R}' \mathcal{P}_{\mathcal{X}} X = \sum_{ij} \kappa_{ij} \langle \mathcal{P}_{\mathcal{X}} X, \omega_{ij} \rangle \mathcal{P}_{\mathcal{X}}(\omega_{ij}) = \sum_{ij} \kappa_{ij} \langle X, \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \rangle \mathcal{P}_{\mathcal{X}}(\omega_{ij}).$$

Namely, $\mathcal{P}_{\mathcal{X}} \mathcal{R}' \mathcal{P}_{\mathcal{X}} = \sum_{ij} \kappa_{ij} \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij})$. Similarly, $\mathcal{P}_{\mathcal{X}} = \sum_{ij} \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij})$. So we obtain

$$\begin{aligned} \|p^{-1} \mathcal{P}_{\mathcal{X}} \mathcal{R}' \mathcal{P}_{\mathcal{X}} - \mathcal{P}_{\mathcal{X}}\| &= \left\| \sum_{ij} (p^{-1} \kappa_{ij} - 1) \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \right\| \\ &\triangleq \left\| \sum_{ij} X_{ij} \right\|, \end{aligned}$$

where $X_{ij} = (p^{-1} \kappa_{ij} - 1) \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij})$ is a zero-mean random variable.

To use Lemma 10, we need to work out M and L therein. Note that

$$\begin{aligned} \|X_{ij}\| &= \|(p^{-1} \kappa_{ij} - 1) \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij})\| \\ &\leq |p^{-1} \kappa_{ij} - 1| \|\mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij})\| \\ &\leq \max\{p^{-1} - 1, 1\} \|\mathcal{P}_{\mathcal{X}}(\omega_{ij})\|_F^2 \\ &\leq \frac{c\mu r}{n_{(2)}p} \\ &\triangleq L. \end{aligned}$$

Furthermore,

$$\begin{aligned}
\left\| \sum_{ij} \mathbb{E} [X_{ij} X_{ij}^*] \right\| &= \left\| \sum_{ij} \mathbb{E} [X_{ij}^* X_{ij}] \right\| \\
&= \left\| \sum_{ij} \mathbb{E} \left(\frac{\kappa_{ij} - p}{p} \right)^2 [\mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij})] [\mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij})] \right\| \\
&= (p^{-1} - 1) \left\| \sum_{ij} \|\mathcal{P}_{\mathcal{X}}(\omega_{ij})\|_F^2 \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \right\| \\
&\leq \frac{c\mu r}{n_{(2)}p} \left\| \sum_{ij} \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \otimes \mathcal{P}_{\mathcal{X}}(\omega_{ij}) \right\| \\
&= \frac{c\mu r}{n_{(2)}p} \|\mathcal{P}_{\mathcal{X}}\| \\
&= \frac{c\mu r}{n_{(2)}p} \triangleq M.
\end{aligned}$$

Since $M/c = 1 > \epsilon$, by Lemma 10, we have

$$\begin{aligned}
&\mathbb{P}\{\|p^{-1}\mathcal{P}_{\mathcal{X}}\mathcal{R}'\mathcal{P}_{\mathcal{X}} - \mathcal{P}_{\mathcal{X}}\| < \epsilon\} \\
&\leq 2mn \exp\left(-\frac{3\epsilon^2}{8M}\right) \\
&= 2mn \exp\left(-\frac{3\epsilon^2 n_{(2)}p}{8c\mu r}\right) \\
&\triangleq 2mn \exp\left(-\frac{C\epsilon^2 n_{(2)}p}{\mu r}\right) \\
&\leq 2mn \exp(-CC_0 \log n_{(1)}) \\
&= 2n^{-CC_0+2},
\end{aligned}$$

where the second inequality holds once we have $p \geq C_0\epsilon^{-2}(\mu r \log n_{(1)})/n_{(2)}$. So the proof is completed. \blacksquare

E. Proofs of Lemma 12

We proceed to prove Lemma 12.

Proof: From the definition of operator R' , we know that

$$\mathcal{R}'(Z) = \sum_{ij \in \mathcal{K}} \langle Z, \omega_{ij} \rangle \omega_{ij} = \sum_{ij} \delta_{ij} \langle Z, \omega_{ij} \rangle \omega_{ij},$$

where δ_{ij} s are i.i.d. Bernoulli variables with parameter p . Notice that $Z \in \tilde{\mathcal{T}}$, so we have

$$Z - p^{-1}\mathcal{P}_{\tilde{\mathcal{T}}}\mathcal{R}'Z = \sum_{ij} (1 - p^{-1}\delta_{ij}) \langle Z, \omega_{ij} \rangle \mathcal{P}_{\tilde{\mathcal{T}}}\omega_{ij},$$

and

$$\langle Z - p^{-1}\mathcal{P}_{\tilde{\mathcal{T}}}\mathcal{R}'Z, \omega_{ab} \rangle = \sum_{ij} (1 - p^{-1}\delta_{ij}) \langle Z, \omega_{ij} \rangle \langle \mathcal{P}_{\tilde{\mathcal{T}}}\omega_{ij}, \omega_{ab} \rangle.$$

We now want to invoke the scalar Bernstein inequality. Let $X_{ij} = (1 - p^{-1}\delta_{ij}) \langle Z, \omega_{ij} \rangle \langle \mathcal{P}_{\tilde{\mathcal{T}}}\omega_{ij}, \omega_{ab} \rangle$ with zero mean.

$$\begin{aligned}
|X_{ij}| &= |(1 - p^{-1}\delta_{ij}) \langle Z, \omega_{ij} \rangle \langle \mathcal{P}_{\tilde{\mathcal{T}}}\omega_{ij}, \omega_{ab} \rangle| \\
&\leq |(1 - p^{-1}\delta_{ij})| \max_{ij} |\langle Z, \omega_{ij} \rangle| \|\mathcal{P}_{\tilde{\mathcal{T}}}\omega_{ij}\|_F \|\mathcal{P}_{\tilde{\mathcal{T}}}\omega_{ab}\|_F \\
&\leq \frac{2\mu r}{n_{(2)}p} \max_{ab} |\langle Z, \omega_{ab} \rangle| \\
&\triangleq L.
\end{aligned}$$

Furthermore,

$$\begin{aligned}
\sum_{ij} \mathbb{E} X_{ij}^2 &= \sum_{ij} \mathbb{E} (1 - p^{-1} \delta_{ij})^2 \langle Z, \omega_{ij} \rangle^2 \langle \mathcal{P}_{\tilde{\tau}} \omega_{ij}, \omega_{ab} \rangle^2 \\
&= (p^{-1} - 1) \sum_{ij} \langle Z, \omega_{ij} \rangle^2 \langle \mathcal{P}_{\tilde{\tau}} \omega_{ij}, \omega_{ab} \rangle^2 \\
&= (p^{-1} - 1) \max_{ij} \langle Z, \omega_{ij} \rangle^2 \sum_{ij} \langle \omega_{ij}, \mathcal{P}_{\tilde{\tau}} \omega_{ab} \rangle^2 \\
&= (p^{-1} - 1) \max_{ij} \langle Z, \omega_{ij} \rangle^2 \|\mathcal{P}_{\tilde{\tau}} \omega_{ab}\|_F^2 \\
&\leq \frac{2\mu r}{n_{(2)}p} \max_{ab} \langle Z, \omega_{ab} \rangle^2 \\
&\triangleq M.
\end{aligned}$$

Since $M/L = \max_{ab} |\langle Z, \omega_{ab} \rangle| > \epsilon \max_{ab} |\langle Z, \omega_{ab} \rangle|$, by scalar Bernstein inequality, we obtain

$$\begin{aligned}
\mathbb{P} \left\{ \max_{ab} |\langle Z - p^{-1} \mathcal{P}_{\tilde{\tau}} \mathcal{R}' Z, \omega_{ab} \rangle| < \epsilon \max_{ab} |\langle Z, \omega_{ab} \rangle| \right\} &\leq 2 \exp \left\{ \frac{-3\epsilon^2 \max_{ab} \langle Z, \omega_{ab} \rangle^2}{8M} \right\} \\
&= 2 \exp \left\{ \frac{-3\epsilon^2 n_{(2)}p}{16\mu r} \right\} \\
&\leq n_{(1)}^{-10},
\end{aligned}$$

provided that $p \geq C_0 \epsilon^{-2} \mu r \log n_{(1)}/n_{(2)}$ for some numerical constant C_0 . ■

F. Proofs of Lemma 13

We are prepared to prove Lemma 13.

Proof: From the definition of operator \mathcal{R}' , we know that

$$\mathcal{R}'(Z) = \sum_{ij \in \mathcal{K}} \langle Z, \omega_{ij} \rangle \omega_{ij} = \sum_{ij} \delta_{ij} \langle Z, \omega_{ij} \rangle \omega_{ij},$$

where δ_{ij} s are i.i.d. Bernoulli variables with parameter p . So

$$Z - p^{-1} \mathcal{R}' Z = \sum_{ij} (1 - p^{-1} \delta_{ij}) \langle Z, \omega_{ij} \rangle \omega_{ij}.$$

Let $X_{ij} = (1 - p^{-1} \delta_{ij}) \langle Z, \omega_{ij} \rangle \omega_{ij}$. To use the matrix Bernstein inequality, we need to bound X_{ij} and its variance. To this end, note that

$$\begin{aligned}
\|X_{ij}\| &= |1 - p^{-1} \delta_{ij}| |\langle Z, \omega_{ij} \rangle| \|\omega_{ij}\| \\
&\leq p^{-1} \|\omega_{ij}\|_F \max_{ij} |\langle Z, \omega_{ij} \rangle| \\
&= p^{-1} \max_{ij} |\langle Z, \omega_{ij} \rangle| \\
&\triangleq L.
\end{aligned}$$

Furthermore,

$$\begin{aligned}
\left\| \sum_{ij} \mathbb{E} X_{ij} X_{ij}^* \right\| &= \left\| \sum_{ij} \mathbb{E} (1 - p^{-1} \delta_{ij})^2 \langle Z, \omega_{ij} \rangle^2 \omega_{ij} \omega_{ij}^* \right\| \\
&\leq p^{-1} \max_{ij} \langle Z, \omega_{ij} \rangle^2 \left\| \sum_{ij} \omega_{ij} \omega_{ij}^* \right\| \\
&= p^{-1} \max_{ij} \langle Z, \omega_{ij} \rangle^2 \|nI_m\| \\
&= \frac{n}{p} \max_{ij} \langle Z, \omega_{ij} \rangle^2.
\end{aligned}$$

Similarly,

$$\left\| \sum_{ij} \mathbb{E} X_{ij}^* X_{ij} \right\| \leq \frac{m}{p} \max_{ij} \langle Z, \omega_{ij} \rangle^2.$$

We now let $M = n_{(1)} \max_{ij} \langle Z, \omega_{ij} \rangle^2 / p$ and set t as $C'_0 \sqrt{p^{-1} n_{(1)} \log n_{(1)} \max_{ij} |\langle Z, \omega_{ij} \rangle|}$. Since $M/L = n_{(1)} \max_{ij} |\langle Z, \omega_{ij} \rangle| > t$, by the matrix Bernstein inequality, we obtain

$$\begin{aligned} \mathbb{P} \left\{ \|Z - p^{-1} \mathcal{R}' Z\| < C'_0 \sqrt{\frac{n_{(1)} \log n_{(1)}}{p} \max_{ij} |\langle Z, \omega_{ij} \rangle|} \right\} &= \mathbb{P} \{ \|Z - p^{-1} \mathcal{R}' Z\| < t \} \\ &= (m+n) \exp \left\{ \frac{-3t^2}{8M} \right\} \\ &\leq n_{(1)}^{-10}. \end{aligned}$$

■

G. Proofs of Lemma 14

We proceed to prove Lemma 14.

Proof: For any fixed matrix Z , it can be seen that

$$\mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}} Z = \sum_{ij \in \Omega_{obs}} \langle \mathcal{P}_{\widehat{\mathcal{V}}} Z, \omega_{ij} \rangle \omega_{ij} = \sum_{ij} \kappa_{ij} \langle Z, \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij} \rangle \omega_{ij}.$$

Note that the operators \mathcal{R}' and $\mathcal{P}_{\mathcal{I}}$ are commutative according to (9), thus we have

$$\mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\mathcal{I}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}} Z = \sum_j \delta_j \sum_i \kappa_{ij} \langle Z, \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij} \rangle \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij}.$$

Similarly, $\mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}} Z = \sum_j \sum_i \kappa_{ij} \langle Z, \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij} \rangle \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij}$, and so

$$(a^{-1} \mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\mathcal{I}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}} - \mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}}) Z = \sum_j (a^{-1} \delta_j - 1) \sum_i \kappa_{ij} \langle Z, \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij} \rangle \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij}.$$

Namely,

$$a^{-1} \mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\mathcal{I}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}} - \mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}} = \sum_j (a^{-1} \delta_j - 1) \sum_i \kappa_{ij} \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij} \otimes \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij}.$$

We now plan to use concentration inequality. Let $X_j \triangleq (a^{-1} \delta_j - 1) \sum_i \kappa_{ij} \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij} \otimes \mathcal{P}_{\widehat{\mathcal{V}}} \omega_{ij}$. Notice that X_j is zero-mean and self-adjoint. Denote the set $g = \{\|C_1\|_F \leq 1, C_2 = \pm C_1\}$. Then we have

$$\|X_j\| = \sup_g \langle C_1, X_j(C_2) \rangle = \sup_g |a^{-1}(\delta_j - a)| \left| \sum_i \kappa_{ij} \langle C_1, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle \langle C_2, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle \right| \triangleq |a^{-1}(\delta_j - a)| \sup_g |f(\delta_j)|.$$

According to (9),

$$\begin{aligned} \|\mathcal{P}_{\widehat{\mathcal{V}}} C_1\|_{2,\infty}^2 &= \max_j \sum_i \langle C_1, \omega_{ij} \widehat{V} \widehat{V}^* \rangle^2 \\ &= \max_j \sum_i \langle G_j^* e_i^* C_1, e_j^* \widehat{V} \widehat{V}^* \rangle^2 \\ &\leq \max_j \sum_i \|e_i^* C_1\|_2^2 \|e_j^* \widehat{V} \widehat{V}^*\|_2^2 \\ &= \max_j \|C_1\|_F^2 \|e_j^* \widehat{V} \widehat{V}^*\|_2^2 \\ &\leq \frac{\mu r}{n}, \end{aligned}$$

where G_j is a unitary matrix. So we have

$$\begin{aligned} |f(\delta_j)| &\leq \sum_i |\langle C_1, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle| |\langle C_2, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle| \\ &= \sum_i \langle C_1, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle^2 \\ &\leq \sum_i \langle \mathcal{P}_{\widehat{\mathcal{V}}} C_1, \omega_{ij} \rangle^2 \\ &\leq \|\mathcal{P}_{\widehat{\mathcal{V}}} C_1\|_{2,\infty}^2 \\ &\leq \frac{\mu r}{n}, \end{aligned}$$

where the first identity holds since $C_2 = \pm C_1$. Thus $\|X_j\| \leq \mu r a^{-1} n^{-1} \triangleq L$. We now bound $\sum_j \|\mathbb{E}_{\delta_j} X_j^2\|$. Observe that

$$\|\mathbb{E}_{\delta_j} X_j^2\| \leq \mathbb{E}_{\delta_j} \|X_j^2\| = \mathbb{E}_{\delta_j} \|X_j\|^2 = \mathbb{E}_{\delta_j} a^{-2} (\delta_j - a)^2 \sup_g f(\delta_j)^2,$$

where the last identity holds because C_1 , C_2 and δ_j are separable. Furthermore,

$$\begin{aligned} \sup_g f(\delta_j)^2 &= \sup_g \left(\sum_i \kappa_{ij} \langle C_1, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle \langle C_2, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle \right)^2 \\ &\leq \sup_g \left(\sum_i |\langle C_1, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle \langle C_2, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle| \right)^2 \\ &= \left(\sum_i \langle C_1, \mathcal{P}_{\widehat{\mathcal{V}}}(\omega_{ij}) \rangle^2 \right)^2 \\ &\leq \left(\sum_i \langle \mathcal{P}_{\widehat{\mathcal{V}}} C_1, \omega_{ij} \rangle^2 \right) \left(\sum_i \langle \mathcal{P}_{\widehat{\mathcal{V}}} C_1, \omega_{ij} \rangle^2 \right) \\ &\leq \|\mathcal{P}_{\widehat{\mathcal{V}}} C_1\|_{2,\infty}^2 \sum_i \langle \mathcal{P}_{\widehat{\mathcal{V}}} C_1, \omega_{ij} \rangle^2 \\ &\leq \frac{\mu r}{n} \sum_i \langle \mathcal{P}_{\widehat{\mathcal{V}}} C_1, \omega_{ij} \rangle^2. \end{aligned}$$

Therefore,

$$\begin{aligned} \sum_j \|\mathbb{E}_{\delta_j} X_j^2\| &\leq \mathbb{E}_{\delta_j} a^{-2} (\delta_j - a)^2 \frac{\mu r}{n} \sum_{ij} \langle \mathcal{P}_{\widehat{\mathcal{V}}} C_1, \omega_{ij} \rangle^2 \\ &= \frac{\mu r (1-a)}{na} \|\mathcal{P}_{\widehat{\mathcal{V}}} C_1\|_F^2 \\ &\leq \frac{\mu r}{na} \\ &\triangleq M. \end{aligned}$$

Since $M/L = 1 > \epsilon$, by the matrix Bernstein inequality,

$$\begin{aligned} \mathbb{P} \left\{ \|a^{-1} \mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\mathcal{I}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}} - \mathcal{P}_{\widehat{\mathcal{V}}} \mathcal{R}' \mathcal{P}_{\widehat{\mathcal{V}}}\| < \epsilon \right\} &= \mathbb{P} \left\{ \left\| \sum_j X_j \right\| < \epsilon \right\} \\ &\leq (m+n) \exp \left\{ \frac{-3\epsilon^2}{8M} \right\} \\ &= (m+n) \exp \left\{ \frac{-3\epsilon^2 na}{8\mu r} \right\} \\ &\leq n_{(1)}^{-10}, \end{aligned}$$

provided that $a \geq C_0 \epsilon^{-2} (\mu r \log n_{(1)})/n$ for some numerical constant $C_0 > 0$. ■

Hongyang Zhang received the Master's degree in computer science from Peking University, Beijing, China in 2015. He is now a Ph.D. student in Machine Learning Department, Carnegie Mellon University, Pittsburgh, USA. His research interests include machine learning, statistics, and numerical optimization.

Zhouchen Lin (M'00-SM'08) received the Ph.D. degree in applied mathematics from Peking University in 2000. He is currently a Professor with the Key Laboratory of Machine Perception, School of Electronics Engineering and Computer Science, Peking University. He is also a Chair Professor with Northeast Normal University. He was a Guest Professor with Shanghai Jiao Tong University, Beijing Jiaotong University, and Southeast University. He was also a Guest Researcher with the Institute of Computing Technology, Chinese Academy of Sciences. His research interests include computer vision, image processing, machine learning, pattern recognition, and numerical optimization. He was an area chair of CVPR 2014, ICCV 2015, NIPS 2015, AAAI 2016, IJCAI 2016, and CVPR 2016. He is an Associate Editor of the IEEE Transactions on Pattern Analysis and Machine Intelligence and the International Journal of Computer Vision.

Chao Zhang (M'06) received the Ph.D. degree in electrical engineering from Beijing Jiaotong University, Beijing, China, in 1995. He was a Post-Doctoral Research Fellow with the National Laboratory on Machine Perception, Peking University, Beijing, from 1995 to 1997. He has been an Associate Professor with the Key Laboratory of Machine Perception, School of Electronics Engineering and Computer Science, Peking University, since 1997. His current research interests include image processing, statistical pattern recognition, and visual recognition.